Using Quantum Resources for Security and Computation

Elliott Mark Ball

Supervised by Prof. Rob J. Young

This thesis is submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

August 2024

Physics Department
Lancaster University

Contents

1 Introduction **15**

2 Background **18**

2.1 Physical Unclonable Functions 18

2.1.1 Optical PUFs 21

2.1.2 Challenge-Response Pair Space 23

2.1.3 Ubiquitous PUFs 25

2.2 Quantum Dot PUFs 25

2.2.1 Semiconductor Devices 26

2.2.2 Computer Vision and Optical Fingerprinting 30

2.3 Quantum Computing & Algorithms 32

2.3.1 Variational Quantum Algorithms in the NISQ Era 32

2.3.2 Preliminaries 34

2.3.3 Distance Measures 35

2.3.4 Quantum Compilation 38

2.3.5 VQA Cost Functions 39

3 Experimental Methods **43**

3.1 PUF Fabrication 44

3.1.1 Single-wavelength emission devices 44

3.1.2 Multiple-wavelength emission devices 44

3.2 Lab-based data capture 45

3.2.1 Shroud-based Capture 45

3.2.2 Microscope-based Capture 46

3.3 Smartphone-based capture 47

3.4 Fingerprinting Algorithms 48

3.4.1 Adapted High Boost 48

3.4.2 Reduced Modified LBP 50

3.5 Processing Pipeline 53

3.5.1 Pre-Processing . . . 53
3.5.2 Post-Processing . . . 55
3.6 PUF Assessment . . . 56
3.6.1 Assessment Model . . . 57
3.6.2 Performance Metrics . . . 59

4 Quantum Dot Physically Unclonable Functions 65
4.0.1 Fingerprinting Outputs . . . 67
4.1 Shroud (Black-box) Conditions . . . 68
4.1.1 Qualitative Comparison . . . 68
4.2 Ambient Lighting . . . 70
4.2.1 Single session comparisons . . . 71
4.2.2 Comparisons against shroud reference . . . 74
4.3 Concluding Remarks . . . 76

5 Hybrid QD PUF 79
5.1 One-to-Many PUF . . . 80
5.2 Correlated Hybrid Extended Entropy Source . . . 81
5.2.1 Correlation as a resource . . . 82
5.3 Shroud, Full PUF Analysis . . . 85
5.3.1 Fingerprinting Outputs . . . 85
5.3.2 Uniqueness and Correlation . . . 87
5.4 Optical Microscope Analysis . . . 90
5.5 Concluding Remarks . . . 92

6 Mixed State Compilation 94
6.1 Quantum Mixed State Compiling . . . 95
6.1.1 QMSC Algorithm . . . 96
6.1.2 Gradient Analysis . . . 98
6.2 Search for a local cost function . . . 100
6.2.1 Local States-Based Attempts . . . 101
6.2.2 Local Measurements on Global States . . . 102

7 Discussion and Future Work 113
7.1 Future Work . . . 114
7.1.1 Authentication . . . 114
7.1.2 Compilation . . . 116

A Towards Implementation **117**

Preface

This thesis is my original work and has not been submitted, in whole or in part, for a degree at this or any other university. Nor does it contain, to the best of my knowledge and belief, any material published or written by another person, except as acknowledged within the text.

For Vimla.

Acknowledgments

"We are convinced that midnight is the first minute of a new day, and we are further convinced that we are entitled to a new day." - Gil Scott-Heron

First and foremost, I wish to extend the utmost thanks to Professor Rob Young, for your excellent support, optimism, and guidance throughout my time at Lancaster. Not only would this thesis not have existed without your help, but thanks to your approach as a supervisor, I have been afforded the freedom to learn so much more on top of what is compiled here.

I must also thank my fellow research group members and those I've shared an office with: James, Kieran, Tom, Sam, Charlie, Edward, Jay, Nema, Ella, Elizabeth, and Fangling. Additional thanks to those at Quantum Base: David, Daniel and Daniel, Hugo, Hector, Blake, and Phil, for your support and useful conversations throughout. I give special thanks to James and Charlotte Fong, for both the support of your friendship, and the many well-needed trips to The Three Mariners. And Angelo, for your help in fabrication of samples; thoughtful discussions on PUF behaviour; but most importantly, an uncountable number of coffee breaks. Special thanks to my collaborators at Los Alamos, particularly Zoë Holmes, for both the opportunity and all that I learned through it.

To all those who were a part of Coulston Road; some for longer, some for shorter, thank you for the warm welcome to Lancaster. Anton Heagney and Ben Clarke especially, for all you put up with, and for those Saturdays. Thank you Amira Zied and Olivia Albrecht, the first friends I made in the city. Thanks to Alex and Gav, as well as the wider Golden Lion community, for providing me with worthwhile distractions and giving me a chance to woo (or bore) the crowds of Lancaster with synthesisers, words, and a broken guitar. Thanks Ben Jackson, the original climbing partner.

To Harry Rosewin, Josh, Matt, Anoosh, and Caitlin: you put up with me throughout undergrad, and, to your own dismay, chose to stick with me thereafter. Thank you all. Special thanks to Christian Porter and Edmund Dable-Heath, for your constructive

conversations throughout both my PhD and Masters studies, as well as the proof-reading and general support and friendship.

I can't not mention the Chuntongo and E11 community, who helped keep me sane during the extended lockdowns, and, to my surprise and great pleasure, have ended up being an instrumental part of my life as well as lifelong friends. Special thanks to Steven, for your useful pointers regarding algebraic decompositions.

To Lucy, who was still helping me with maths problems at postgraduate level despite stopping at undergrad; and Molly, thank you for keeping me in check. And to Bradley, Denna and Aimee. Kyle, the Gang: cheers to ACO! Harry J. Cockerell, thanks for your insanity and being the most loveable idiot. M, whose support rings ever dear despite the thousands of miles between us, thank you.

Thank you to my family, for your ever-lasting support throughout these four years. Despite going down a journey that bemuses you all, your support has never wavered. Mum, despite what comes, the combination of strength and love that you display without fault has forever inspired me, and without having learnt those two dear qualities from you and Dad, four years of devotion would have been impossible. Luke and Zara, you might be my younger siblings, but I still learn from you each day. To my older siblings, Michael, Serena, and Charlotte; thank you for your love and support. Serena, not only have you supported me directly with your love and big sisterly advice; but you've given me even more through my niece and nephew, Annabelle Lucy and Max James. No matter how dark a day may be, our family will always supply a form of light, with Annie, Max, Luke, Zara, and Christopher leading the procession.

And to Madee: you have supported me in more ways than could possibly be counted. Your kindness, your care, and your constant willingness to lend a helping hand have done much more for me than I could ever fit in a single page (or two).

"You shouldn't let poets lie to you." - Björk

Contributions

Fabrication of the ink used in the samples discussed in chapters 4 and 5, as well as the deposition of said ink, was performed by Angelo Lamantia at Lancaster University. Code for the iOS app used to control capturing conditions for the experiments discussed in chapter 4, as well as the appendix, was written by Daniel Abreu, in conjunction with the author, at Lancaster University. Classical simulations of the quantum algorithm discussed throughout chapter 6 were performed by Nic Ezzell and Aliza U. Siddiqui, at Los Alamos National Laboratory. Implementations of the algorithm on IBM superconducting devices was performed by Nic Ezzell at Los Alamos National Laboratory.

List of Publications

Elliott M. Ball, Kieran D. Longmate, Joonas Majaniemi, Angelo Lamantia, Daniel Abreu, Matthew J. Fong, Ella Mann-Andrews, and Robert J. Young. *"Smartphone-based Fingerprint Extraction from Quantum-Optical PUFs"*, Under preparation, 2023.
—The work in this manuscript is discussed in great detail, and expanded upon, in chapter 4. All measurements in manuscript were performed by the author, with analysis led by the author, as well as writing of the manuscript.
Kieran D. Longmate, **Elliott M. Ball**, Edmund Dable-Heath, Robert J. Young. *"Signing information in the Quantum Era"*, AVS Quantum Sci. 2, 044101, 2020.
—Author contributed to the writing throughout this manuscript, and led the research and writing of section II.
Kieran D. Longmate, Nema M. Abdelazim, **Elliott M. Ball**, Joonas Majaniemi and Robert J. Young . *"Improving the longevity of optically-read quantum dot physical unclonable functions"*, Scientific Reports volume 11, 10999, 2021.
—The work in this manuscript is of relevance to the useability of the devices discussed in chapters 4 and 5 of this thesis. The author contributed to development of algorithms used in analysis, as well as analysis itself.
Matthew J. Fong, Christopher S. Woodhead, Nema M. Abdelazim, Daniel C. Abreu, Angelo Lamantia, **Elliott M. Ball**, Kieran Longmate, David Howarth, Benjamin J. Robinson, Phillip Speed, and Robert J. Young. *"Using intrinsic properties of quantum dots to provide additional security when uniquely identifying devices"*, Scientific Reports volume 12, 16919, 2022.
—The work in this manuscript relates to potential secondary layers of security for optical PUF devices. The author contributed to the research and analysis presented.
Nic Ezzell, **Elliott M. Ball**, Aliza U. Siddiqui, Mark M. Wilde, Andrew T. Sornborger, Patrick J. Coles, and Zoë Holmes, *"Quantum mixed state compiling"*, Quantum Sci. Technol. 8 035001, 2023. —This manuscript details the principal algorithm discussed in chapter 6, and provides numerical simulations of the algorithm's performance, and experimental data. The author contributed to the formulation of the algorithm, analysis of the cost functions used in optimisation, and the writing of the manuscript.

Abstract

Quantum mechanics and information theory have jointly impacted multiple fields. Two in particular are security and computing. Via the use of quantum resources, exploits in currently used digital security systems are known, whilst the theory also promises security for future systems. Quantum theory has been shown to have fundamental impacts on computing technology, but modern experimental hardware is limited in power and use cases.

This thesis is concerned with developments in the use of quantum resources in both fields. Physically unclonable functions (PUFs), a static form of entropy source with uses in hardware-based cryptography, are investigated. Utilising colloidal quantum dot based ink in order to fabricate a series of optical PUF (OPUF) devices, the reliable transformation of (classical) optical information whose source's fundamental optical properties are governed by quantum theory into a unique fingerprint for further processing in cryptographic protocols is explored. First, the ability to use only a smartphone device to both excite, and capture the optical emission of, an OPUF is explored. It is shown that these images can be reliably converted into binary keys via two algorithms. Next, a novel type of OPUF is proposed. Two inks, each comprised of quantum dots with peak emission at different wavelengths are used to fabricated a device which produces two, separable responses under a single optical challenge. The correlation between two outputs from a given device is found to be inconsistent, with the cause for such inconsistencies explored.

Finally, by making use of a hybrid quantum-classical computing method, an algorithm for learning the preparation circuit of an unknown mixed state is defined. In order to combat known issues with scalability of current hardware, this work explores the possibility of reformulating the well-known Hilbert-Schmidt distance using local quantum objects. A variety of functions are investigated, with the final answer remaining open.

List of Figures

2.1 Schematic modelling a PUF as blackbox process. 20
2.2 Schematic modelling of an arbitrary dynamic entropy source as a blackbox process. 21
2.3 Schematic demonstrating the instrumental idea behind an optical PUF. . 22
2.4 Visual representation of the relationship between challenges and responses for a given PUF device. 24
2.5 Figure demonstrating the difference in bandstructures 26
2.6 Simple band diagrams showing examples for the: creation of free electron (a), creation of hole (b), creation of electron-hole pair (c) 27
2.7 Density of states for semiconductors with different degrees of freedom. . . 29
2.8 Schematic of electron-hole creation and re-combination. 30
2.9 A simple representation of the principle behind variational quantum algorithms. 33
2.10 Visual representation of Bloch Sphere formalism. 35
2.11 Quantum circuit representing the most generic quantum operation. . . . 36
2.12 Example of Loschmidt-Echo type circuit. 40
2.13 Simple visual representation of cost function landscapes exhibiting different phenomena. 42

3.1 Flowchart displaying a simplefied view of the overall, generic processing pipeline for interrogating PUF devices in this work. 44
3.2 QD-PUF design. 45
3.3 Simple schematic showcasing the shroud-based capture system. 47
3.4 Simple schematic of the microscope-based capturing system. 48
3.5 Visual schematic demonstrating experimental set up for capturing images of PUF devices using a smartphone. 49
3.6 Example of AHB algorithm on a given input. 50
3.7 Example of R-MLBP algorithm on a given input. 51
3.8 Simple representation of pixel-wise comparison points for RMLBP. 52

3.9 An example of the intra and inter Hamming distance distributions generated for a given PUF instance. 60

4.1 Examples of R-MLBP and AHB outputs for three different PUF input images. 67
4.2 Example intra-Hamming and inter-Hamming distributions for fingerprints produced across a range of QD PUFs under R-MLBP and AHB. 68
4.3 Plot of calculated average biases for PUF fingerprint outputs. 69
4.4 Decidability scores for all 48 single-emission PUF devices. 70
4.5 A comparison of false acceptance rate and expected number of independent bits for a single-dot emission PUF device, imaged in shroud. 71
4.6 A comparison of false acceptance rate and expected number of independent bits for a single-dot emission PUF device, 0.06 kLux. 72
4.7 A comparison of false acceptance rate and expected number of independent bits for a single-dot emission PUF device, 0.326. 73
4.8 A comparison of false acceptance rate and expected number of independent bits for a single-dot emission PUF device, 0.220 kLux. 74
4.9 Intra-Hamming and inter-Hamming distributions generated via R-MLBP from images captured at 0.006 kLux, 0.326 kLux, and 0.220 kLux 75
4.10 Example of the average decidability calculated for a PUF device, over varying radius for fingerprints generated via R-MLBP and AHB. 76
4.11 4-panel figure of thresholded FARs and FRRs. 77

5.1 Diagramatic representations of one-to-one and one-to-many PUFs. 80
5.2 Expected emission spectra for absorbance and emission for CIS530 and CIS650 colloidal quantum dots. 82
5.3 Schematic showing potential cross-talk between quantum dots in HPUFs. 83
5.4 Schematic showcasing potential usecase for HPUF devices with independent correlation maps. 84
5.5 Filtered images of DH-PUF and WH-PUF devices. 86
5.6 Examples of captures emission, and output R-MLBP and AHB fingerprints for H-PUF devices. 87
5.7 Figures showcasing intra- and inter- distributions over varying choice of radius for fingeprints generated from images of isolated emissions for a given H-PUF device. 88
5.8 Variation of DOF and bias for a given HPUF device. 89

5.9 Examples of the varying statistics for Hamming distributions over choices of radius for R-MLBP and AHB for isolated CIS650 emission for a given HPUF device. 90
5.10 Optical microscope images of DHPUF and WHPUF devices. 91
5.11 Annotated isolated emission patterns of CIS5530 and CIS650 HPUF devices. 92

6.1 Overview of QMSC 97
6.2 QMSC Subroutines 98

A.1 Simple schematic to prompt a smartphone user to interrogate a PUF device.117
A.2 Example inputs and outputs into a perspective correction algorithm. . . . 118

Chapter 1

Introduction

Since the mid-20th century, how we view the processing of information has changed significantly. From a practical standpoint, the ever-growing ubqiuitous presence of personal computers in the home, office, and on the person, (initially a luxury reserved for large businesses, with the financial and spatial resources to house large machines for what is now viewed as elementary data processing) have further changed how we perceive and interact with information.

From a more theoretic standpoint, Claude Shannon's work in the 1940s helped create a basis of understanding information through the lens of mathematics. Whilst earlier work by Nyquist [88] and Hartley [56] had examined elements of the physical nature of information transmission (the latter's name is given to the unit of information in base 10 due to his work on information distinguishability), Shannon's work is considered a historical turning point in the development of information theory. Via discussions with von Neumann, Shannon was made aware of the direct link between his mathematical theory of communication and earlier work in statistical and thermodynamics, resulting in the formal definition of entropy in an informational sense, which is considered to encompass precisely the same notion that researchers had paired it with earlier. This idea that information theory could be linked to physical theories pervaded and strengthened, and following Landauer's infamous proclamation in 1991 that "Information is physical" [71], has been taken as gospel by the majority of physicists.

The 20th century also bore witness to the modern development of quantum theory, and the realisation that the inherent workings of nature may (in a simplefied view) boil down to chance. If we are to subscribe to Landauer's viewpoint and accept a direct relationship between (phyiscal) nature and the concept of information, it is then natural to wonder how the advent of quantum mechanics reshapes our understanding of information theory. Indeed, in 1962, Gordon [49] first considered the effect of quantum phenomena on the transmission of information through a channel. Despite this question

not being considered prior to Gordon's work, a quantum analogue of Shannon's entropy had already been defined by von Neumann in his work on statistical quantum mechanics [113], and a multitude of alternative measures of information in quantum theory have since been introduced [96, 108]. Feynman's proposition in 1982 [43] that the simulation of quantum systems would likely be impossible via classical devices builds further on the link between information processing and physical theories, and, along with Deutsch's proposal of the quantum equivalent to a Turing machine [30], helped birth the study of quantum computing. Since, quantum computers have been realised, and the field of quantum information theory has given rise to an increased understanding of cryptography and (more generally) digital security, via the (both theoretical and experimental) processing of quantum information within physical systems.

The use of quantum theory to achieve security goes beyond the direct manipulation of quantum information. In the combat against counterfeit goods, the study of Physical Unclonable Function (PUFs) has gained much traction. A PUF is a physical system that exhibits some unique and readable property when interrogated, due to randomness inherent within the system. Their production of (seemingly) random bits of information has led to them being dubbed an *entropy source*, in a similar vein to random number generators [18]. Devices that are capable of generating high levels of entropy are of great practical use in the field of cryptography, where it's generally held that the higher the randomness of information (typically keys in this context), the greater the level of security. Multiple types of PUFs have been proposed and experimentally realised, typically exploiting seemingly random artefacts caused by hard to control manufacturing processes of a particular device, such as microscale variation in paper fibres [115]. As well as exploiting classical physical processes, we are able to realise *quantum* PUFs, exploiting the non-deterministic[1] nature of quantum theory in order to harness the unique and unclonable response. By documenting and classifying how a PUF responds differently to different interrogation parameters allows for their use as a means of secure authentication: if it is trusted that a PUF device cannot be genuinely cloned, then it is sufficient to verify these challenge and response pairs in order to determine the authenticity of the device. Use of PUFs in an underlying protocol has many applications across both physical security, and potential cryptographic tasks [101, 79].

In this work, we focus on how information can be manipulated and processed in light of quantum theory, to aid the fields of security and computation. Chapter 2 seeks to

[1]Of course, the non-determinism of quantum theory is not a closed question, and the field of quantum foundations seeks to understand precisely how quantum theory describes our reality. However, for the purposes of this work, whether or not the universe truly is deterministic is inconsequential, as we have no way to exploit any potential superdeterminism of the universe. As such, we are safe to assume non-determinism, and leave the philosophical nature of such questions to the foundational theorists.

establish a basis from which the research work is contextualised, via an exploration of related work and the theoretical background behind the topics. Chapter 3 then provides a systematic account of the experimental methods with which results were both obtained and analysed.

Chapters 4-5 focus on the use of quantum dots, semi-conductor nanomaterials whose optical interactions are governed by quantum mechanics, in order to build secure authentication protcols. An examination of the classical optical, unique information produced by these devices as a result of optical interactions governed by quantum mechanics is conducted, in order to build secure authentication protcols. In chapter 4, that these materials' optical properties can be captured at a macroscale is investigated as means for conducting authentication schemes using them with ubiquitous technology, namely, a smartphone. Chapter 5, shifting away from the ubiquitious technology framework, proposes a novel device utilising two layers of ink comprising of quantum dots, for further research in applications to authentication.

Chapter 6 deals with quantum information directly: focusing on how we can use quantum resources in order to understand and process quantum information, in the form of (sets of) two-dimensional quantum states. A novel algorithm for the compilation of arbitrary mixed states is presented, with the focal work being the construction and analyses of candidate cost functions for use in the algorithm.

Chapter 2

Background

This chapter provides a brief review of work relevant to the remainder of the thesis, and explores the theory and background of the work shown in the later chapters. Firstly, previous work relating to physically unclonable functions (PUFs) is explored, with particular focus given to those of an optical and/or ubiquitous nature, as well as the first algorithm proposed for fingerprinting PUFs. Then, an overview of early work in quantum computing and algorithms is provided, with an explicit focus on variational quantum algorithms (VQAs), and early work on state compilation.

Firstly, the field of Physical Unclonable Functions (PUFs) is introduced, along with the key notion of challenge response pairs, and examples of the latter's relation to the former. A small recap of the physical background of quantum dot PUFs is provided, before giving a brief introduction to computer vision techniques, and their relevance to optical PUFs.

Then, variational quantum algorithms are introduced, via an introduction to current-term quantum computing and its challenges. Preliminaries for the work presented in chapter 6, along with mathematical notation used throughout said chapter, are given here. The task of mixed state compilation, which forms the basis of the work presented throughout the aforementioned chapter, is then introduced, before an exploration of scalability issues for the algorithm presented in this work, and VQAs in general.

2.1 Physical Unclonable Functions

A form of hardware-based security, PUFs have been proposed as a primitive for a variety of cryptographic schemes, notably key generation and authentication [106]. PUF devices are categorised by a *challenge-response mechanism*; given a (set of) "challenge(s)", i.e., some physical interrogation of the device, some subsequent and readable (set of) response(s) is recorded. As sources of entropy, each challenge-response mechanism appears random

— both when compared to other challeng-response pairs (CRPs) from the same device, and when compared to CRPs from other devices. As such, it is expected that any PUF (generally) satisfies the following requirements:

- Uniqueness — No two PUFs should exhibit the same readable property.
- Stability — Given a PUF device, the property to be read should provide the same information over a long period of time.
- Assymetric fabrication process — A PUF device should be easy to fabricate, but comparatively hard to replicate.
- Unpredictability — It should be hard to predict the response for a given challenge, even if other CRP outputs are known.
- Scrambling — Joint with unpredictability, the process by which a PUF converts a challenge into a response should be hard to discover.

Depending on the implementation of a PUF device, some of these requirements may be loosened. Where appropriate, these will be addressed later.

The need for stability inherently characterises PUFs as *static* entropy sources, that is to say, they produce random information within the spatial domain (between different devices, and under different inputs), but their outputs are static over time: the output for a given input at time t_0 should be the same at time $t_1 > t_0$, as shown in fig. 2.1. This is in direct contrast to *dynamic* entropy sources (shown in fig. 2.2 as a contrast), such as (true) random number generators (RNGs) [18], which must produce different random information, even under the same input, at times t_0 and t_1. Dynamic entropy sources can act as a primitive for a variety of cryptographic purposes, such as key generation, and (when true randomness is guaranteed) are essential for the implementation of one-time pads, and thus provably unbreakable encryption. Static entropy sources may also be used for key generation, but such uses will usually be paired with conditions on implementation, such as an assumption that the bearer of the entropy source is a particular party. This makes them natural candidates for use for authentication (which will be a driving motivation in this work), whilst their ability to produce a reliable, unclonable uniqueness makes them suitable for identification and anti-counterfeiting purposes (such purposes need not be distinct from authentication, as explored in chapters 4 and 5).

PUFs were first proposed by Pappu, under the moniker "Physical One-Way Functions" [90], taking inspiration from the cryptographic tool of (mathematical) one way functions [62]. This early PUF was proposed in the form of a three-dimensional inhomogeneous microstructure consisting of micron-scale glass spheres cured in optical-grade

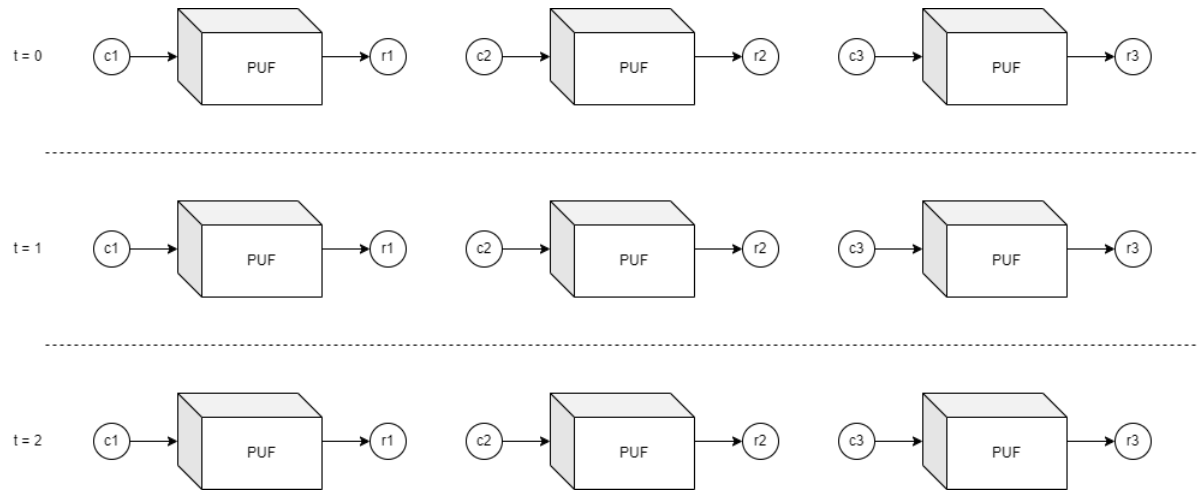

Figure 2.1: Schematic modelling a PUF as blackbox process. At time $t = 0$, under different challenges (represented by $c1, c2, c3$), the PUF produces different responses, (represented by $r1, r2, r3$). At subsequent times $t = 1$ and $t = 2$, the PUF produces the same response when re-interrogated with a previous challenge.

epoxy. Taking advantage of the optical coherent scattering nature of such structures, these tokens were interrogated using a laser beam, with the resulting optical speckle pattern being captured with a camera, and treated as the response. In order to extract a binary key from the output speckle pattern, the first fingerprint extraction algorithm for PUFs was also proposed, relying on the use of Gabor Transformations to characterise the different optical patterns produced by different challenges (formed by rotations of the incident beam at different angles) and devices. With both the challenge and subsequent response consisting of optical information, PUFs such as these are typically dubbed *Optical PUFs (O-PUFs)*. In the following year, Gassend et. al [47] coined the term PUF when introducing their *Silicon Physical Random Function*. With the aim of characterising, identifying, and authenticating semiconductor devices on integrated circuits, these devices take an (electrical) signal as an input, and a measurement of signal delay is taken as the output, forming the challenge-response mechanism for an *all-electronic PUF*. Similarly to the original (optical) PUF proposal, a variety of challenge-response pairs can be assessed for each device, with different input signals (acting here as the challenge) leading to a different output delay measurement. Since, an increasing amount of research has led to a large variety of electronic PUFs [47, 73, 59, 97, 60, 85], which may be characterised further depending on the intrinsic physical system governing the electrical challenge-response mechanism [84], with potential input challenge signals for just electronic PUFs ranging from input voltage, to input binary vectors, the latter used by arbiter PUFs to interrogate different internal delay mechanisms, resulting in a variety of multiplexer outputs as responses [114] .

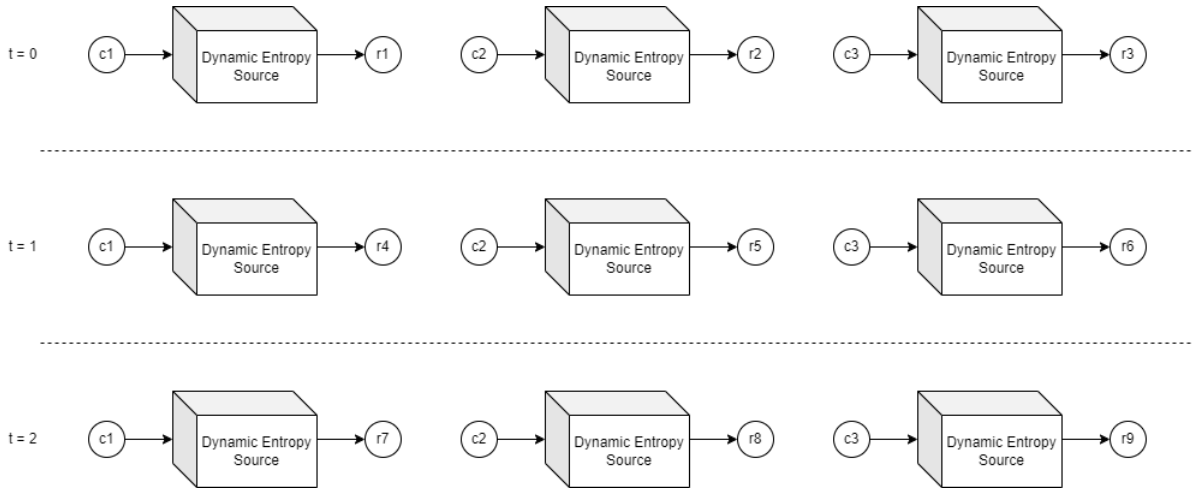

Figure 2.2: Schematic modelling of an arbitrary dynamic entropy source as a blackbox process. At time $t = 0$, under different challenges (represented by $c1, c2, c3$), the entropy source produces different responses, (represented by $r1, r2, r3$). At subsequent times $t = 1$ and $t = 2$, even when presented with the same set of challenges, different responses are produced.

As described above, a single PUF device typically produces multiple challenge-response mechanisms. By considering the relationship between PUF size and the number of accessible (typically, differentiable) Challenge-Response pairs (CRPs), PUFs of all types may instead be characterised by this relationship. Typically, devices whose number of CRPs scales well with device size are considered *strong PUFs*, whilst those whose number of CRPs scales in a relatively limited manner are considered *weak PUFs* [99]. Challenge response spaces are explored further in subsection 2.1.2, and chapter 5.

The work presented in chapters 4 and 5 deals specifically with a form of optical PUF. For a full description of the different categories (and subcategories) of proposed PUFs, the reader is directed to [84].

2.1.1 Optical PUFs

As described above, the original optical PUF proposal relies on optical scattering by an inhomogenous microsctructure. However, a variety of optical PUFs whose randomness stems from different intrinsic processes have since been proposed. Here, we will provide a brief run-through of some of those considered thus far.

The use of scattering materials has proven popular in PUF fabrication since Pappu's initial proposal. For information stored on compact disc format, their have been a variety of proposals for authentication by considering the scattering behaviour of the disc itself when interrogated by lasers [54], typically due to random variations in the manufacturing process. Exploiting Rayleigh backscattering in the transmission of information via fibre

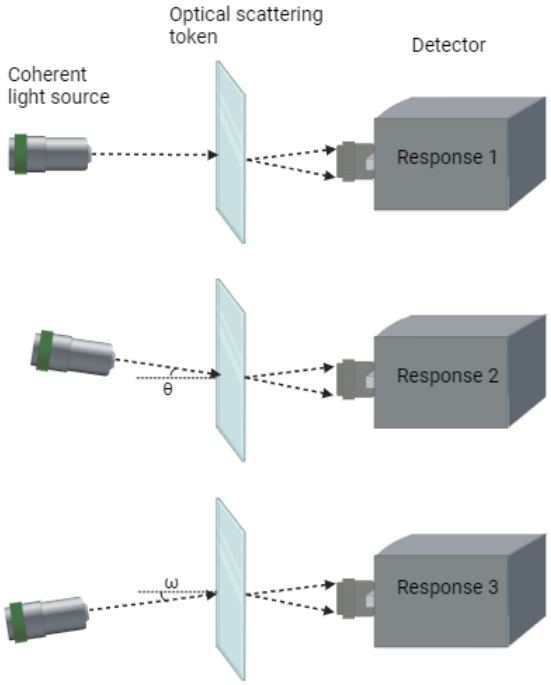

Figure 2.3: Schematic demonstrating the instrumental idea behind an optical PUF, consisting of a scattering token. The PUF is interrogated using a coherent light source, and following some interaction between the incident light and PUF, transmitted light is captured by a detector to interpret the challenge. The changing angle of incidence of light (relative to the plain; top, 0, middle, θ, bottom, ω) represents different challenges, for which the detector will record distinct responses.[1]

optical cables has been proposed as a method to verify the authenticity of information transfers, due to uncontrollable variations in the manufacturing of the fibre. Recently, 'visually hidden' porous nanostructure layers have been exploited for their scattering properties [70], producing OPUFs whose microstructure cannot be easily probed (for example by electron microscopy), further obfuscating the processes yielding a unique response for enhanced security.

Silicon-based PUFs have been introduced as an alternative method of intrinsically tying an optical PUF process to the fabrication of electronic devices [52, 74]. A related area, that of optical PUFs based on interferometers, has also been of interest [64]. Another subset of OPUFs that has gained traction is the *paper PUF*, in which characteristics of paper fibres (which are inherently random due to uncontrollable manufacturing features) are used to uniquely identify a device [9]. These typically rely on gathering information on the reflection of incident light, as opposed to the complex scattering in microstructures exploited in Pappu's case. In 2015, Wong and Wu proposed a method for feature detection in paper, achievable using smartphone cameras [115]. Paper PUFs are

[1]Parts of this figure were created with Biorender.com

typically proposed for applications associated with banknotes: that they can characterise an intrinsic property of the paper used is a key benefit, allowing for the use of a PUF without introducing a new manufacturing step for the product.

All OPUFs discussed so far utilise a fixed CRP space, that is, once a PUF device has been manufactured, all available CRPs are determined and unchangeable. Recently, Gan et. al. proposed a *reconfigurable* PUF [46], utilising the phase transition of VO_2 nanocrystals in order to alter the laser speckle pattern of the device, allowing for an expansion of the CRP space post-fabrication. Expansion of the CRP space had previously been considered via the combination of different challenges, such as the influence of optical interference discussed in [98].

When considering the security and use cases for PUFs, the intrinsic physical processes that produce the unique responses are also considered. This can be easily divided into those whose uniqueness stems from classical mechanics; such as those that take advantage of stochastic, hard to predict fabrication processes; and those whose uniqueness stems from quantum mechanics; such as those which make use of the probabilistic nature of quantum tunnelling, and/or semiconducting structures. The use of classical processes leads to a potential weakness in the future: should technology advance to the point where the manufacturing process can be much more finely controlled, an adversarial party may be able to examine a device sufficiently to reproduce the minute defects that control its uniqeuness, and thus obtain a clone of the device. On the other hand, devices which harness the power of quantum mechanics should remain hard to replicate. The fabrication of defect-free monolayers is understood to be difficult, and thus the ability to precisely replicate a device's bandgap structure should also remain difficult. Other forms of quantum PUFs have been proposed which make direct use of quantum states [104], and duplication of such devices is strongly ruled out via the no-cloning theorem. Devices which tie the uniqueness of the quantum material with a hard-to-predict manufacturing process (similar to the case of classical optical PUFs) have been exhibited. One such example is PUFs consisting of semiconducting quantum dots [63], which gain an advantage due to the ability to finely control emission properties via the fabrication of dots of different sizes, and the added complexity of how dot molecules interact with each other when forming clusters. Such materials have been shown to have a further strength, with their unique (non-linear) relation between excitation strength and emission acting as an additional layer of security [44].

2.1.2 Challenge-Response Pair Space

The relationship between an input challenge and output response is pivotal to the working, and understanding of PUF devices. As discussed previously, the density of challenge-

response pairs (CRPs) for a device is an important aspect in characterising *weak PUFs* and *strong PUFs*. These characterisations further inform us on the type of application that we can expect a given PUF to be useful for: typically, strong PUFs are considered good candidates for interactive challenge-response authentication protocols, whilst weak PUFs are typically viewed as good candidates for the basis of key generation protocols. However, given the nascency of the field, these extended characterisations (high density implying authentication, and low density implying key generation) should not be viewed as universal: depending on the specifics of the protocol, even weak PUFs may be extended into authentication devices, which is a major focus of the work presented in chapters 4 and 5.

The mathematical analogy of one-way functions that gave PUFs their name may be extended to the relationship between a given challenge and response; with the mapping typically being one-to-one, i.e. for a given challenge, c, there typically exists only one unique response, r, to which c is mapped. This is diagramatically shown in figure 2.4. Formally, we may write

$$f : \mathcal{C} \to \mathcal{R}; f(x) \mapsto y \tag{2.1}$$

where, for a given challenge, x in the domain $\mathcal{C}$, f may be considered as the PUF device mapping x to its unique element, y of the codomain, $\mathcal{R}$.

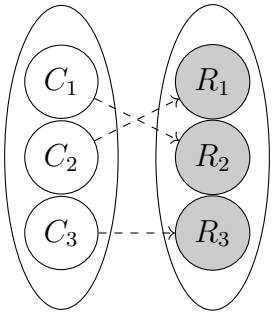

Figure 2.4: Visual representation of the relationship between challenges (left) and responses (right) for a given PUF device. Analogously to injective functions, for each element in the domain, the mapping yields one unique element of the codomain.

The work presented in this thesis falls strictly under the realm of weak PUFs, with applications in authentication with ubiquitous technology, where the same optical response will need to be obtained in varying conditions. The work in chapter 4 deals only with a single challenge, and single response, whilst in chapter 5, discussions of CRP spaces are expanded, showcasing a PUF for which a single challenge elicits two responses.

2.1.3 Ubiquitous PUFs

One of the major blocks in the widescale adoption of optical PUFs by the average person is accessibility. The vast majority of PUF propositions require specialised knowledge and equipment, such as laboratory-grade optics and microscopes in order to correctly interrogate and analyse the challenge-response mechanism. One method to increase the usability of PUFs is to remove the need for direct interaction between the user and the PUF. For electronic PUFs, adaptation to specific use cases allows for frameworks in which widespread use may be adopted soon. Proposed frameworks typically link an electronic PUF to the *'Internet of Things'* [101], such as those proposed for smart vehicle authentication [7, 5].

Recent work on PUFs that do require user interaction has sought to remove technological barriers, such as the need for specialised equipment. In 2019, Liu et. al. [75] proposed a quantum dot-based inkjet-printed device that may be interrogated using a smartphone and a USB-connected microscope, paving the way for easy-to-interrogate PUF devices. More work utilising smartphone-based microscopy has been conducted since, with Zhang et. al. [118] proposing a multimodal optical PUF fabricated via growth of diamond microparticles on heterogeneous structures, which may be interrogated using a smartphone microscope. They also propose further layers of security for device authentication, via the fingerprinting of "low level" optical information imaged simply, and "higher level" optical information captured via photoluminescence measurements and dark-field scattering recordings. Similarly to the earlier mentioned work from Gan et. al., Zhang et. al. also propose that their PUF is dynamic and reconfigurable, in this case via air oxidation, and showcase that different fingerprints may be reliably captured from the same device following reconfiguration.

2.2 Quantum Dot PUFs

The use of colloidal solutions to manipulate travelling light has a long history. Since the Medieval period, the distribution of small particles within a glass matrix has been used for artistic and religious endeavours [36, 57], with stained glass windows featuring in British churches since the 7th century, whose transmission of 'heavenly light' [72] influenced the school of Gothic architecture through to the 12th century and beyond [28]. Despite the prevalence of colloidal solutions throughout society and architecture, it was not until the 20th century that the scientific basis for such optical properties was probed, contributing to the development of optical nanocrystal theory. For a look at the history of colloidal solutions in the development of nanocrystal theory, see [36]. For an understanding of how (specifially) colloidal quantum dots exhibit unique optical properties, the next

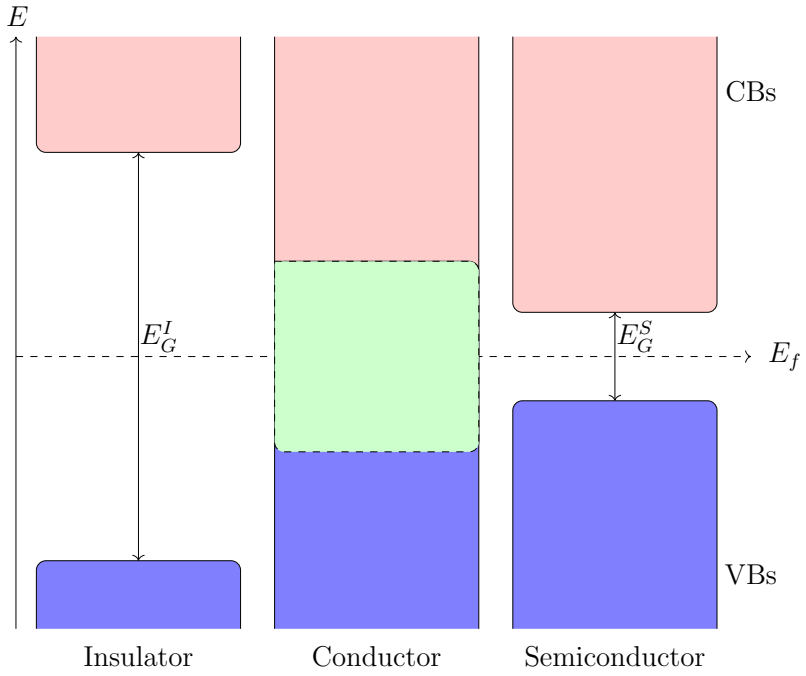

Figure 2.5: Figure demonstrating the difference in bandstructuresbetween (left to right) Insulators, Conductors, Semiconductors, with conduction bands (valence bands) represented by CBs (VBs). Here, E_f represents the Fermi level, and $E_G^{I(S)}$ represents the gap energy for insulators (semiconductors), whilst the green overlap represents the band overlap in conductors.

part of this chapter will focus on the theory of semiconductors, before shifting towards performance metrics for PUFs, and how we will choose to consider them.

2.2.1 Semiconductor Devices

The results presented in chapters 4 and 5 of this thesis exploit the optical properties of semiconducting nanostructures, specifically quantum dots, for the creation of unique digital fingerprints. Here, a brief theoretical background on semiconducting devices will be presented, highlighting the mechanisms that allow for the creation of unique optical patterns, and subsequently, their fingerprints.

2.2.1.1 Bulk materials

Before considering the impact of spatial confinement on semiconducting crystals of small sizes, we begin by introducing the (general) mechanisms by which bulk materials may act as semiconductors. The conductivity of a material may be characterised by its *band structure*, i.e., the energy levels which are permissible physical states for an electron and the levels which are not permissible, the *bandgaps* (a visual representation of these

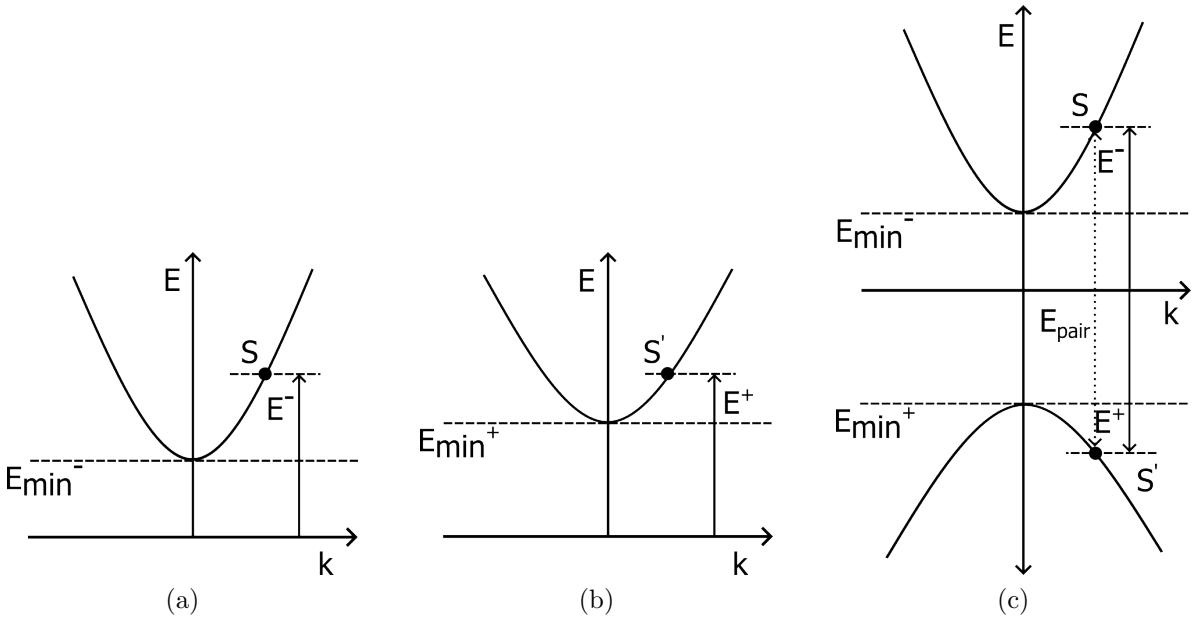

Figure 2.6: Simple band diagrams showing examples for the: creation of free electron (a), creation of hole (b), creation of electron-hole pair (c); with energy plotted against momentum. E^{-}_{min} and E^{+}_{min} represent the minimum energy required for the creation of a free electron and hole, respectively.

is shown in figure 2.5). Specifically, the energy levels (*bands*) that exist closest to the Fermi level (the highest energy level an electron can occupy at 0K) of the material are of most interest, with these being the *valence* and *conduction bands*. For an insulator, the bandgap is large, and spans over the Fermi level, necessitating a large change in energy for the transfer of charge carriers from the valence band to the conduction band. In a conductor however, there is an overlap between the two bands (with the Fermi level existing within this overlap), allowing for free movement of electrons and conductivity of electricity. For semiconductors, the bandgap is small, allowing electrons to be moved from the valence band to the conduction band with only a small amount of energy (with a dependence on temperature).

When an electron-hole pair is created in bulk materials, typically the coulomb attraction between the electron and hole allows for the correlation of their motion, and they can subsequently be treated as a quasiparticle, namely, the *exciton*. Such excitons may move freely through the semiconductor, modelled as a particle in a (large) box, with the permissible energy bands and unconfined k-space allowing for a continuous distribution of states (as seen in the first panel of Fig. 2.7).

2.2.1.2 Small Semiconductor crystals

So far, we have considered how bulk semiconductors behave, with excitons typically being treated as free particles. As quantum dots are specifically crystals of semiconducting material whose size sits somewhere between that of a single molecule, and bulk crystalline structures, the size of the structure, as well as the effects of confinement, must be considered. As confinement is introduced, the model changes. With 2D semiconductors (i.e., confined in one dimension), the model becomes that of a potential well, whilst for 1D semiconductors the model becomes that of a quantum wire. For semiconducting crystals with (sufficiently) small radius, R, it ceases to be suitable to treat the exciton as a free particle. We re-adopt the particle in a box model, whilst taking the confinement in all dimensions into account. Examining the time-independent Schrödinger equation

$$\frac{-\hbar^2}{2m}\nabla^2\Psi = E\Psi, \tag{2.2}$$

we have the boundary conditions for the box model, $\Psi(0) = \Psi(L) = 0$, which yield the solutions

$$E_n = \frac{\hbar^2}{2m_{\text{ex}}}\left(\frac{\pi n}{L}\right)^2 \tag{2.3}$$

where m_{ex} is the effective mass of the exciton, and L is the confinement scale. As such, at small confinement scales (approximately that of the de Broglie wavelength), the allowed energy levels cease to be continuous, and are instead quantised with the density of states being (roughly) represented by a delta function (with slight variations in breadth of the delta function peak, according to the precise size of the QD). Figure 2.7 shows the density of states as a function of energy for each type of confined semiconductor, showcasing how this density tends towards discrete values as degrees of freedom is reduced.

2.2.1.3 Optical Properties

In this work, quantum dots were chosen as a building block for optical PUFs over other candidates due to their unique optical properties, particularly the fine-tunable wavelength of peak emission, and high quantum yield. These properties are a result of their underlying semiconductor structure. In quantum dots, the decaying of an exciton (i.e., subsequent recombination of exciton into an electron in the valence band) results in a radiative release of energy in the form of a photon, whose energy is determined by the bandgap properties of the semiconducting nanocrystal. Therefore, by controlling the size of the nanocrystal (and subsequently controlling the size of the bandgap), it is possible to fine-tune the peak emission for these devices. From the afore discussed particle in a box model, the spacing of allowed energy levels widens as the radius decreases. The shifts in

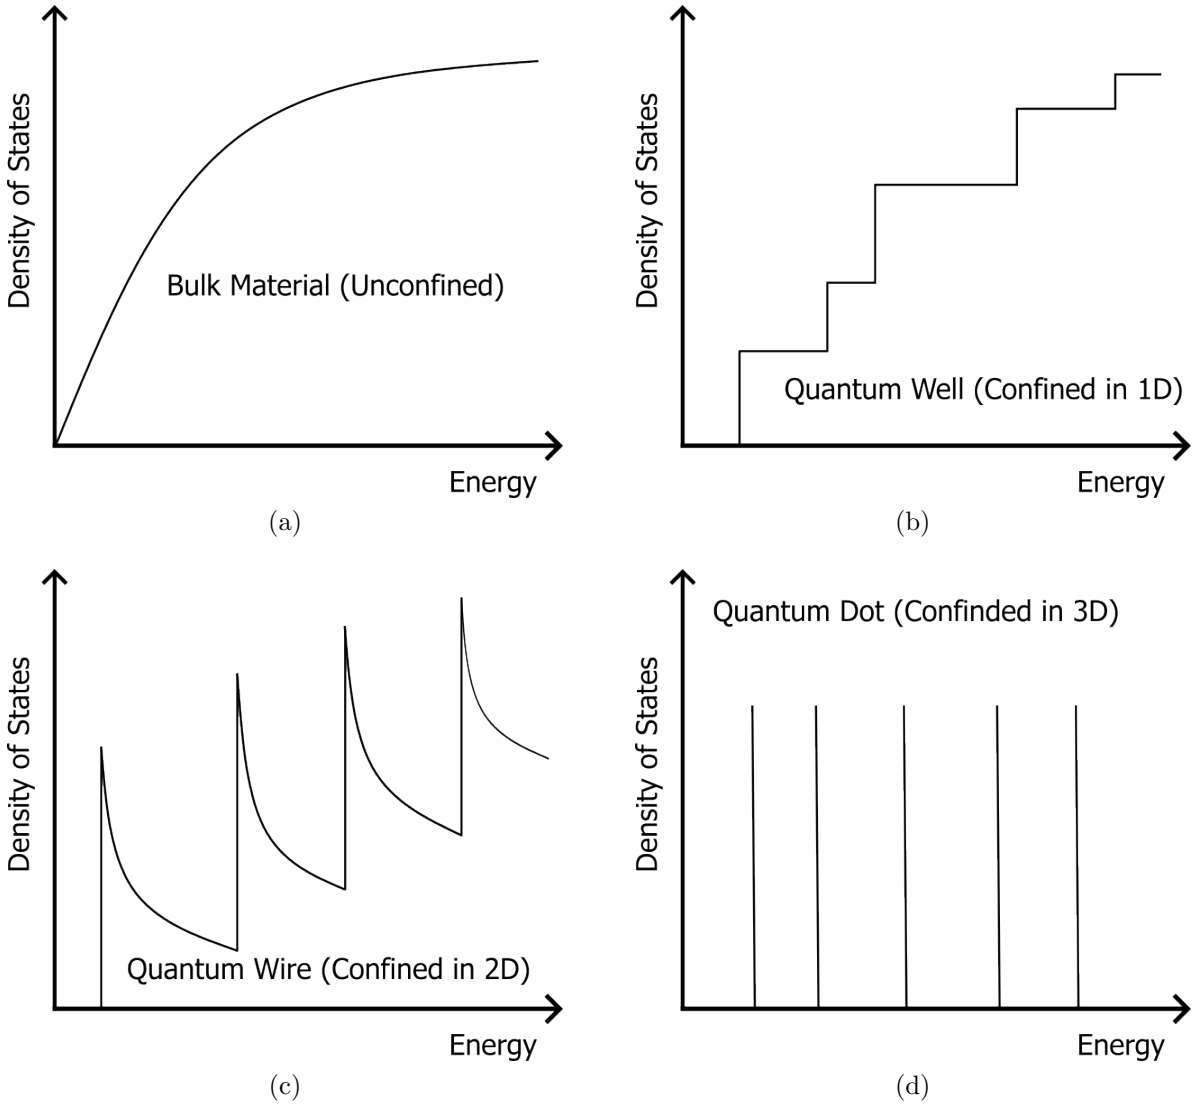

Figure 2.7: Density of states for semiconductors with different degrees of freedom. When unconfined (i.e., three degreeso of freedom) there is a continuous distribution of available states (a). For confinement in one dimension, the available states take the form a piece-wise step function (b), tending to a piece-wise set of peaks for confinement in two dimensions (c). When confinement occurs in all directions, the available states are roughly equivalent to discrete values represented by a delta-function (d).

emission wavelength may be expressed as a function of the dot's core diameter, r:

$$E = E_G + \frac{\hbar^2}{8m_{\text{ex}}}\left(\frac{\pi}{r}\right)^2, \tag{2.4}$$

allowing for the fine-tuning of a dot's optical properties.

However, semiconducting devices often allow for non-radiative mechanisms of recombination, as well as radiative ones. Where both radiative and non-radiative mechanisms

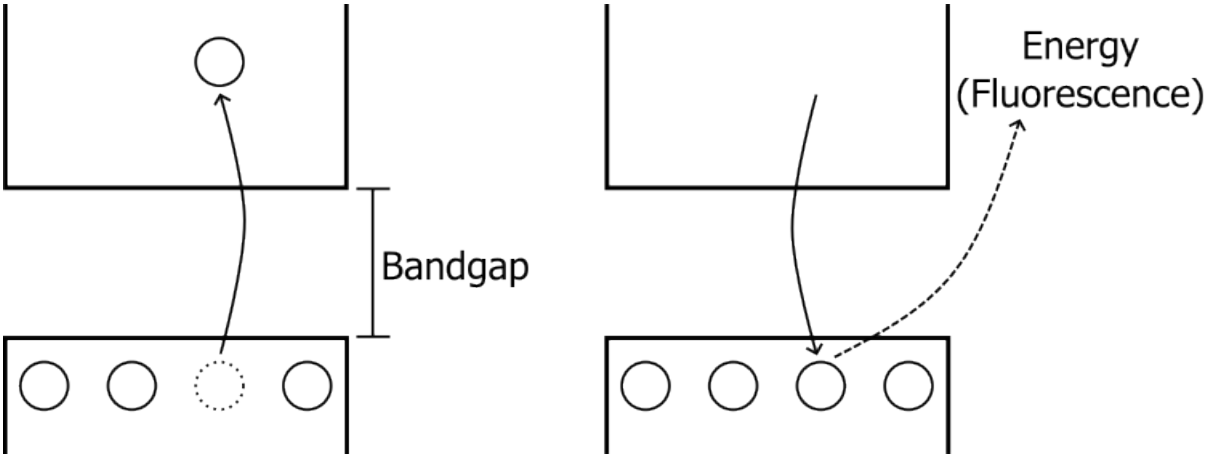

Figure 2.8: Left: Schematic of an electron transferring from the valence band of a semiconuctor device to the conduction band, leaving behind a hole. Right: Schematic of electron and hole recombining, with their system returning to the ground state and outputting energy in the form of fluorescence.

of recombination are present, the quantum yield of the subject material is determined by the recombination lifetime of both mechanisms;

$$\eta_Q = \frac{\frac{1}{\tau_{\text{Rad}}}}{\frac{1}{\tau_{\text{Rad}}} + \frac{1}{\tau_{\text{Non}}}}, \tag{2.5}$$

where τ_{Rad} (τ_{Non}) is the radiative (non-radiative) recombination lifetime. Due to imperfections in manufacturing of semi-conductor devices, further pathways are often opened up for non-radiative mechanisms of recombination, such as Auger or Shockley-Read-Hall recombination. In direct bandgap semiconductors, the non-radiative recombination lifetime is typically much longer than that of the radiative lifetime, resulting in a higher quantum yield when compared to indirect bandgap semiconductors.

The use of CuInS/ZnS core/shell quantum dots allows for the exploitation of both of these strong optical properties: the resulting bandgap created when forming core/shell quantum dots is easily tuneable due to the relative ease of controlling the core and shell sizes, whilst their direct bandgap structure ensures a high quantum yield, and a resulting strong photoluminescence.

2.2.2 Computer Vision and Optical Fingerprinting

The processing of the captured optical information from a PUF into a unique key (that may be further processed for cryptographic protocols) relies on having a method to reliably extract the uniqueness of the optical pattern, and condensing this information into a (binarised) key that retains said uniqueness. Away from PUF research, algorithms re-

lated to such activities have been proposed and researched, mainly within the computer science community, typically to characterise and/or identify structural information within images. The field of computer vision encompasses such algorithms, and has been a topic of research since the 1960s, shortly after the coining of the term artificial intelligence, in an effort to mimic elements of human vision, and give computers an understanding of an image scene. Elements of computer vision often have an overlap with the field of image processing, including tasks such as edge detection. Improvements in computer vision, mainly feature detection algorithms, naturally go hand in hand with advancements in machine learning and artificial intelligence, with the automation of feature recognition providing advancements in fields ranging from physical security [86, 3] to medical imaging and diagnostics [77, 68], with such technology reported as outperforming humans [35].

2.2.2.1 Gabor Hash

In the same work that first introduced the notion of a PUF [90], the Gabor Hash algorithm was also first defined, as a means for extracting a binary fingerprint from the output of a speckle-pattern optical PUF. The algorithm makes use of the Gabor transform [45], a special case of the short-time Fourier transform (STFT), formed by multiplying the function to be transformed by the standard Gaussian function (here, acting as the necessary window function for an STFT), and subsequently transforming via the standard Fourier transform. The Gabor transform is typically used in time-frequency analysis, and is formally defined as follows;

Definition 2.6 (Standard Gabor transform of a signal)**.**

$$G_x(\tau, \omega) = \int_{-\infty}^{\infty} x(t) e^{-\pi(t-\tau)^2} e^{-i\omega t} dt, \tag{2.7}$$

where $x(t)$ denotes the signal to be transformed; and τ and ω constitute the time and frequency domains respectively.

The academical history of Pappu's proposed algorithm serves as an excellent showcase of the ways in which information-theoretic ideas from different fields may inform one another. As Daugman illustrated [27] when proposing the 2D generalisation of the Gabor transform, Gabor's initial (1D) work was instrumental in the development of ideas pertaining to how visual information is processed physiologically. Daugman's work on expanding this understanding to 2D spatial profiles furthered the understanding of simple cells in the primary visual cortex, factoring in the processing of 2D orientation information, and paved the way for Pappu's hash algorithm, which makes use of this 2D generalisation. For a full overview of the implementation and theory of the algorithm,

the reader may consult the original thesis work [90] or the subsequent (summary) report [91].

2.3 Quantum Computing & Algorithms

In 1984, not long after the proposition of quantum machines for information processing by Feynman; Bennett and Brassard proposed the seminal BB84 [17] scheme for quantum key distribution. A year later, Deutsch famously envisioned the *universal quantum computer* [31]. Since Deutsch's description of the universal quantum computer, interest in both the development of quantum algorithms to outcompete classical methods, and the development of hardware to run such algorithms, has grown significantly. Just three years later, Yamamoto and Igeta proposed the first physical realisation of a quantum computer [61], and within another three years, Ekert exploited entanglement as a resource [37] to propose an alternative to Bennet and Brassard's key distribution scheme. In 1992, Deutsch and Jozsa published one of the first deterministic quantum algorithms that could solve a problem exponentially faster than a classical alternative [32]. With these developments, and work done [13, 103] showcasing (mathematically) quantum parralelism [87, p. 30], it began to be widely understood that quantum information processing, and the development of quantum computing, would form an important aspect of future computer science endeavours. In 1996, Seth Lloyd's work [76] demonstrated that Feynman's conjecture regarding the simulation of quantum systems via quantum devices to be correct, adding further interest to the experimental realisation of such devices. Within a year, Cory et. al. [23], and independently, Gershenfeld and Chuang [48], had showcased how NMR computing could be used to realise quantum logic gates, with such a system being experimentally realised for a working 2-qubit device capable of solving quantum algorithms [65, 21, 20] in 1998.

2.3.1 Variational Quantum Algorithms in the NISQ Era

Since the introduction of experimentally realisable NMR based quantum computation, research in quantum hardware has gained a great deal of pace, resulting in somewhat of a race to build machines that feature increasing numbers of qubits, whilst building on work to retain coherence in such machines. This has brought us up to the current state of quantum computing, the so-called 'Noisy Intermediate-Scale quantum era', or *NISQ* computing. The term NISQ, as coined by Preskill [93], characterises the current frontier of quantum computers, whose size is limited to around 50-a few hundred qubits (intermediate-scale), and whose gate sequences introduce a (non-trivial, limiting) amount

of noise over the computations. This contrasts the ideal case, in which devices are highly scalable to a (relatively) large amount of qubits, and whose computations are fault-tolerant. In the last few years, an emerging idea is that NISQ devices, despite their limitations, will still prove useful for certain tasks, so long as algorithms constructed for them are designed in such a way as to mitigate some of these shortcomings. This has led to the nascent field of *Variational Quantum Algorithms*, (VQAs), for which the principal idea is that, by constructing your problem such that it is solved by finding optimal parameters for a parametrised circuit, operations performed on a quantum device can be made less resource-intensive, by offloading other tasks to a classical device. Typically, this might look like running a parametrised quantum circuit and performing some measurement on the NISQ device, whilst performing optimisation (e.g., via gradient descent) on a classical device. This method of utilising both quantum and classical devices to solve a VQA problem is often termed a *hybrid quantum-classical* approach to problem-solving. Fig. 2.9 provides a simplefied visual framework of this idea.

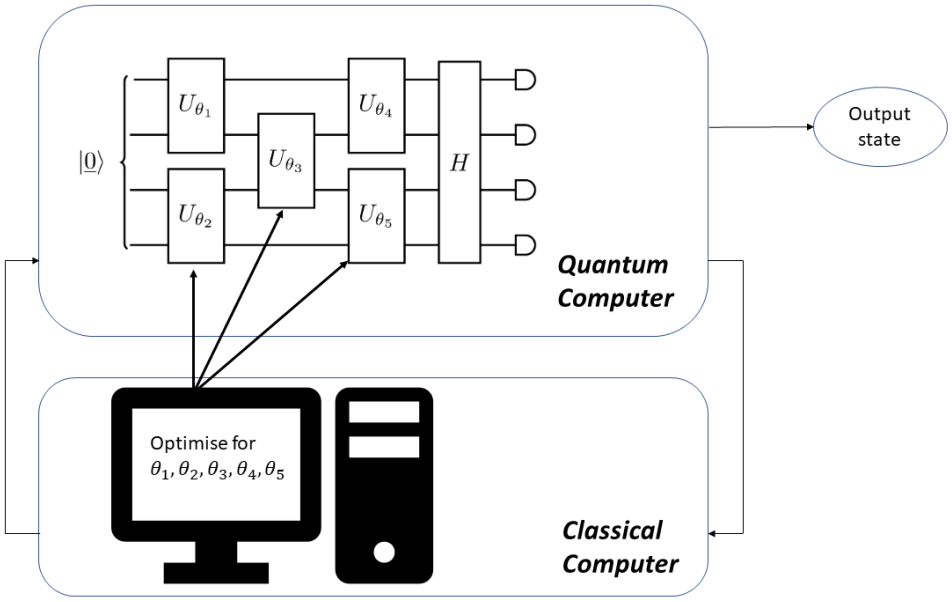

Figure 2.9: A simple representation of the principle behind variational quantum algorithms, via the lens of VQE. A quantum circuit that takes some input (here, the all 0 state), and prepares a trial state via a set of parametrised unitaries, with a known Hamiltonian H, before measurement, is shown. This circuit can be ran on NISQ hardware (represented by the upper box). Measurement results are passed on to a classical computer for the purposes of cost function optimisation. After optimising for each θ_i, the updated parameters are input back into the circuit, and this process is iterated over.

This approach was first introduced by Peruzzo et. al. [92], via the *Variational Quantum Eigensolver* (VQE) algorithm. In essence, VQE utilises parametrised quantum circuitry to prepare a system on a quantum device that models the physics of some

sought after (potentially physical) trial wavefunction, $|\psi(\vec{\theta})\rangle = U(\vec{\theta})\,|\psi\rangle_{in}$. Here, $U(\vec{\theta})$ is the unitary equivalent to the quantum circuitry, parametrised by $\vec{\theta}$, and $|\psi\rangle_{in}$ is some input trial state for the algorithm, typically the all-zero state $|\vec{0}\rangle$. Given a (known) Hamiltonian, H, that represents the system that we wish to compute eigenvectors and eigenvalues for, one can perform a measurement equivalent to

$$\langle\psi|_{in}\, U^{\dagger}(\vec{\theta}) H U(\vec{\theta})\,|\psi\rangle_{in}\,, \tag{2.8}$$

which is lower-bounded by the ground state energy of the Hamiltonian, i.e.,

$$E_0 \leq \frac{\langle\psi(\vec{\theta})|\,H\,|\psi(\vec{\theta})\rangle}{\langle\psi(\vec{\theta})|\psi(\vec{\theta})\rangle}. \tag{2.9}$$

So, by minimising the result of this measurement (such that it approximates the ground-state energy E_0), one can find the relevant eigenvector and eigenvalue within the information encoded in the trial state $|\psi(\vec{0})\rangle$. The VQE algorithm has been shown to have applications in chemistry, but is still limited in the system sizes it can efficiently solve for [42].

This methodology of 'evaluate on quantum, optimise on classical' has been employed for various algorithms since, providing efficient solutions for problems such as time-evolution simulation [82, 38, 12], classical optimisation [4], and quantum machine learning [14]. The strengths and pitfalls of the hybrid VQA approach are explored in section 2.3.5.

2.3.2 Preliminaries

2.3.2.1 Quantum Information and Tensor Algebra

This thesis will utilise standard bra-ket notation for representing qubits, in which the traditional basis elements are represented as $|0\rangle = \begin{pmatrix}1\\0\end{pmatrix}$ and $|1\rangle = \begin{pmatrix}0\\1\end{pmatrix}$. Arbitrary mixed states are represented via their density operators, e.g., $\rho = |\psi\rangle\langle\psi| \in \mathcal{H}$, and this density matrix picture (often associated with the Bloch-sphere formalism, pictured in fig. 2.10), will be heavily utilised throughout chapter 6. When dealing with multi-qubit computations, each of the n qubits resides in its own Hilbert Space, such that the entire system of qubits may be represented as the tensor product between each qubit, residing in the tensor product of each Hilbert Space, i.e., $|x_1\rangle\langle x_1|_{\mathcal{H}_1} \otimes |x_2\rangle\langle x_2|_{\mathcal{H}_2} \otimes \cdots \otimes |x_n\rangle\langle x_n|_{\mathcal{H}_n} \in \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots \otimes \mathcal{H}_n$. For shorthand, when it is unambiguous to do so, we may write such systems as $|x_1x_2...x_n\rangle\langle x_1x_2...x_n|$.

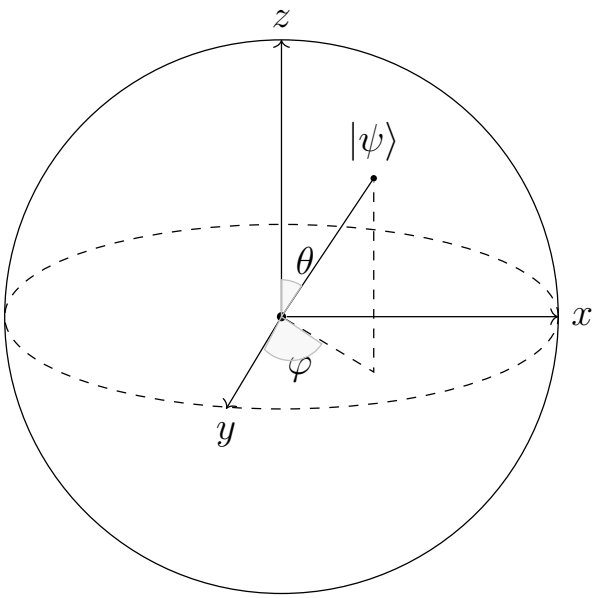

Figure 2.10: Visual representation of Bloch Sphere formalism of quantum mechanics, with a pure state, $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$, labelled. Geometrically, pure states are all those that occupy the surfaces of the Bloch sphere, whilst mixed states occupy the remaining volume. The particular case of the maximally mixed state occupies the centre of the Bloch sphere.

2.3.2.2 Quantum Circuitry

For the purposes of illustrating algorithmic implementation on quantum devices, we also present a formalism of quantum information processing within the framework of quantum circuits. A quantum circuit consists of quantum gates, represented by unitaries, acting on qubits (represented by individual wires) from left to right, typically culminating in a measurement. Fig. 2.11 shows a diagramatic representation of a single qubit quantum operation in the most generic terms via Stinespring's dilation theorem [105]: the system qubit (uppermost qubit here) and an ancillary system (the lower qubit here) is introduced, some quantum operation (the gate, $U^\dagger$) acts on the system and ancillary registers, before the system qubit is measured and the ancillary system is discarded (i.e., traced out). The construction of general operations viewed through this lens is sometimes referred to as *going to the Church of the Large Hilbert Space*, following John Smolin's remark [50].

2.3.3 Distance Measures

In order to assess the similarity between two quantum objects (where here, an object could be a state, a circuit, a measurement, etc.), a form of *distance measure* is necessary. Typically, a used distance measure will be mathematically equivalent to a *metric*.

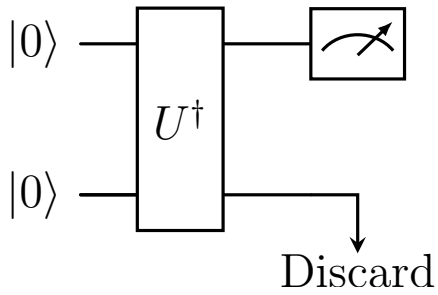

Figure 2.11: Quantum circuit representing the most generic quantum operation, in which any map can be represented by the introduction of an additional system, a unitary, and a discarding of this second (ancillary) system.

Definition 2.10. Given a space S and a function $d : S \times S \to \mathbb{R}$, call d a metric if it satisfies the following:

1. *Non-negativity* — $d(a, b) \geq 0 \ \forall \ a, b \in \mathcal{S}$ (with equality iff $a = b$).
2. *Symmetry* — $d(a, b) = d(b, a) \ \forall \ a, b \in \mathcal{S}$.
3. *Triangle inequality* — It holds that $d(a, c) \leq d(a, b) + d(b, c) \ \forall \ a, b, c \in S$.

Example 2.11. It is worth noting that the *Hamming distance*, used extensively in the analysis presented in chapters 4 and 5 is formally considered a distance measure, where the Hamming distance between two strings (of equal length) is the number of substitutions required to make the strings equal. I.e., in the use of the Roman alphabet, the Hamming distance between ***Belt*** and **Ball** would be 2, whilst the Hamming distance between two binary strings, say, $x = 11101010, y = 01101101$ would be 4.

Mathematical metrics are typically induced by a *norm*, which can be thought of as way to establish the distance between a point and some agreed upon zero point (origin), i.e., the magnitude of an object $a \in S$, with respect to the norm.

Definition 2.12. Given a space S and a function $\| \cdot \| : \mathcal{S} \to \mathbb{R}$, call $\| \cdot \|$ a norm if it satisfies the following:

1. *Non-negativity* — $\|a\| \geq 0 \ \forall \ a \in \mathcal{S}$, with equality iff $a = 0$.
2. *Absolute homogeneity* — $\|\alpha a\| = |\alpha| \|a\| \ \forall \ a \in S$ and any $\alpha \in \mathbb{R}$.
3. *Triangle inequality* — $\|a + b\| \leq \|a\| + \|b\| \ \forall \ a, b \in S$.

Norms are often associated with vector spaces, and thus the term *vector norm* is often used interchangeably with norm throughout the literature. Here, we will be principally concerned with norms on operators, and as such will not typically associate norm with vector norms.

Definition 2.13. For a real number $p \geq 1$ and a vector space V, the p-norm of an n-dimensional vector $\vec{x} \in V$ is defined by

$$\|\vec{x}\|_p = \left(\sum_i^n |\vec{x}_i|^p \right)^{1/p}. \tag{2.14}$$

A choice of $p = 2$ produces the Euclidean norm, whilst a choice of $p = 1$ produces the *taxicab,* or, *Manhattan norm.*

Norms need not only be defined for vectors, but can also be defined on operators (amongst other mathematical objects). For instance, there exists the well known *operator norm*[2], which, for a matrix A, is equivalent to the square root of the largest eigenvalue of $A^\dagger A$. Working in Hilbert spaces, there exists a class of operator norms similar to the p-norms on vectors, dubbed the *Schatten p-norms*, which will be of relevance to this thesis.

Definition 2.15. Given a bounded operator between two Hilbert spaces, $A : \mathcal{H} \to \mathcal{H}$, and a real number $p \geq 1$, the Schatten p-norm of A is defined by

$$\|A\|_p = \left(\mathrm{Tr}\left[|\mathrm{A}|^{\mathrm{p}} \right] \right)^{1/p}. \tag{2.16}$$

The aforementioned operator norm is equivalent to the Schatten ∞-norm, whilst a choice of $p = 1$ yields the *trace class norm*, and $p = 2$ emits the *Hilbert-Schmidt norm*, which will be used in the construction of cost functions in Ch. 6. Due to both the (vector) p-norm and the Schatten p-norm sharing similar names, and the same notation, throughout literature they are often only distinguished via context. Throughout this thesis, any use of the notation $\|\cdot\|_p$ will refer to the Schatten p-norm (unless otherwise specified), and the term p-norm will not be used to refer to the Schatten p-norm, to avoid confusion.

2.3.3.1 Distances between quantum objects

The act of distinguishing between two quantum objects, $|\psi_x\rangle$ and $|\psi_y\rangle$ is a fundamental one throughout quantum computing and information. The simplest way to attempt state discrimination is by measurement, but it is known that measurement can only distinguish non-orthogonal states with a certain probability. Instead, we could calculate the difference between them by some metric on their Hilbert spaces. Shifting to the Bloch-Sphere framework of discussing quantum states (writing $\rho = |\psi_x\rangle\langle\psi_x|$ and $\sigma = |\psi_y\rangle\langle\psi_y|$), it is

[2]The conventional naming of which is unfortunate, obscuring the use of operator norm for an arbitrary norm on operators. In this thesis, the term operator norm will be used in the more general sense, whilst $\|\cdot\|_\infty$ will denote *the* operator norm.

clear that the aforementioned Schatten p-norms are of use for calculating the similarity between two quantum objects, and thus distinguishing between them.

There are also physically meaningful methods by which the similarity of two states, which do not satisfy the requirements necessary to be considered a metric. One well known measure of closeness between two states is the *Fidelity*, defined as

$$F(\rho, \sigma) = \left(\mathrm{Tr} \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right)^2 . \tag{2.17}$$

Whilst not being a formal metric, the fidelity of two states is a widely adopted means of assessing their similarity due to its physical relevance. This relevance is easy to see in the pure case state, where we have that the fidelity simplifies to

$$F\left(\rho, \sigma\right) = |\left\langle \psi_\rho | \psi_\sigma \right\rangle|^2, \tag{2.18}$$

the overlap between the two states. Uhlmann's theorem [110] allows us to extend this to the mixed state case, where we get that the fidelity is equivalent to the maximum overlap between purifications of ρ and σ. *Further, the fidelity induces the* Bures metric*, a quantum-equivalent of the Fisher information metric.*

2.3.4 Quantum Compilation

Compilation of a quantum state is a primitive in quantum computing. With the experimental realisation of quantum computers, it is natural to question how, given some (potentially unknown) quantum state ρ, one can prepare that quantum state on hardware. Compilation can also serve as a primitive for other tasks, or in some cases, perform other tasks as a byproduct. In the case of ρ being unknown, learning its preparation (i.e., learning a unitary for compilation) is evidently equivalent to learning the state itself, which forms one of the most universal tasks accomplished by compilation. Further examples of tasks achievable via compilation will be explored in Chapter 6. Work on compilation is, naturally, not limited to that of states. Since Barenco et. al.'s [8] and Deutsch et. al.'s [33] work on elementary and universal logic gates for quantum computation, early compilation work [109, 55] focused on translating theoretical quantum algorithms into sets of such gates and their unitary matrix representations, for processing on hardware. As work on experimental realisations of quantum hardware progressed, compilation work focusing on specific hardware instances, such as those based on Ising couplings [16], also picked up.

In 2008, Maslov et. al. [81] investigated the recompilation of known circuits into hardware specific instances with a great reduction in depth and gate number. In 2018,

Venturelli et. al. [112] related the task of native gate compilation to the AI field of temporal planning to realise a circuit compilation scheme for circuits with a high degree of commutability, specifically quantum alternating operator ansatz circuits [41, 53]. Later that year [15], this work was extended with the investigation of approaches based on *constraint programming*, in which a task is formulated as a scheduling problem. It was shown that, by exploring a hybrid approach consisting of both temporal planning and constraint programming, one could compile circuits with greater success and scaleability than with either constraint programming or temporal planning alone. More and more, researchers began to make use of machine-learning techniques in the pursuit of compilation. Cincio et. al. presented an ML-based approach to learning the (quantum) algorithm for state overlap [22], which can be viewed as equivalent to (but not limited to) state compilation.

Compilation via variational algorithms has recently been explored too, albeit primarily for pure state compilation. The aforementioned work in [22] teased the variational approach proposed in [69]: an algorithm for pure state compilation that employs the *Hilbert-Schmidt Test* as a cost function is proposed, whilst also introducing a localised variant of the cost function for mitigating the problems invoked by barren plateaus (both local cost functions and barren plateaus will be formally defined and explored as concepts in section 2.3.5 and chapter 6). Sharma et. al. [102] explore the inherent resilience to noise in such approaches, finding that, for the tested circuit types, the output gate sequence of a variational compilation algorithm typically reduces noise (as opposed to inheriting, and perhaps increasing) when compared to an already known gate sequence. Further work by Jones and Benjamin [66] explores compilation for specific input states, via eigensolving techniques and energy minimisation. It is shown that such an approach is robust to gate noise when seeking near-perfect compilation.

2.3.5 VQA Cost Functions

As briefly discussed in chapter 2, variational quantum algorithms have emerged as a hot topic in research, given their strengths in task-oriented programming for current low-qubit devices where purely quantum error correction is hard to achieve. Performance of a VQA has been shown to be dependent on the cost function used for evaluating performance (and hence, "closeness" to the sought-after solution), and the construction of the ansatz used in preparing a trial input to the algorithm. In the case of cost functions, researchers typically seek to ensure that any such function used is

- Faithful — The cost function should vanish if and only if the output of the quantum algorithm matches the desired target solution.
- Efficiently computable on quantum hardware — Evaluation of the cost function

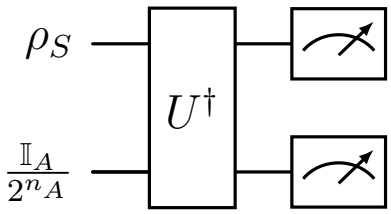

Figure 2.12: Example of Loschmidt-Echo type circuit to measure $\mathrm{Tr}[\rho\sigma]$. Here, $U^\dagger$ aims to "echo" the unitary required to have evolved ρ_S from an initial state.

should be achievable with (relatively) low gate-depth and qubit resources.

- Scaleable — As the problem size grows, complexity of the task should grow linearly with it.

- Operationally meaningful — Non-zero values should provide some physical or operational description about the objects in question.

For problems directly related to physical systems, such as VQE, the formulation of a cost function comes naturally: by computing the energy of the trial state, $|\psi\rangle$, the operational meaning is trivial, as is faithfulness (i.e., if searching for the ground state, the cost function will only be minimised for the ground state). Efficiency naturally depends on the complexity of the Hamiltonian in question, but it is at least known that the Hamiltonian may be decomposed into a series of unitaries that can be adapted and chosen to complement the elementary gate alphabet of a given device. For some other tasks, where the goal is abstracted away from a natural, physical system; such as compiling, compression, and foundational questions such as VQA explorations of consistent histories; defining a suitable cost function is less intuitive, particularly in terms of operational meaning. Cases such as compression and compilation, where "closeness" of states may be measured by fidelity, will allow for an intuitive operational meaning, but less efficient computation. Meanwhile, the use of efficiently computable circuits such as so-called "Loschmidt echo"-style circuits, (named as such as they seek to mimic the effect of a Loschmidt Echo [94, 51] (the composition of a time-reverse unitary evolution and the original unitary evolution)) (Pictured in fig. 2.12) allow for calculations of forms of $p-$norms, introduced in definition 2.13, which can be tied to an operational meaning via known bounds on such norms. However, regardless of the ease of cost function construction, scaleability remains an issue amongst VQAs.

2.3.5.1 Barren Plateaus

Whilst the field of VQAs has shown a great deal of promise in advancing our ability to perform a variety of tasks on quantum hardware, such algorithms typically suffer from the *curse of dimensionality*; as the size of a problem scales up (and hence, the number of qubits used), and as gate depth increases, issues with trainability arise. One phenomenon that has been shown to contribute to this so-called 'curse' is that of *barren plateaus* (BPs). Put simply, a barren plateau is a feature of a cost landscape, in which the cost function gradient is vanishing at the majority of points in the landscape, and therefore, optimisation is made difficult.

Early research in to barren plateaus [83] had shown that they were linked to deep circuits, especially with random initialisations of ansätze. Cerezo et. al. [19] then went on to show the prevalence of BPs even in shallow parametrised circuits, showing that cost functions defined in terms of global observables lead to exponentially vanishing gradients in terms of the number of qubits, n; i.e., barren plateaus appear quickly. However, they also showed that, for shallow circuits, if the cost function is instead formulated in terms of local observables, gradients vanish at worst polynomially in n, and thus the curse of barren plateaus may be staved off for well-devised problems. The difference between these two types of cost functions (global and local) is as follows: global cost functions revolve around global observables, and thus typically involve a final measurement on all qubits. Local cost functions however, instead make use of a series of measurements on $k < n$ qubits, with early research simply focused on the case that $k = 1$. Formally, for a quantum circuit parametrised by some vector $\vec{\theta}$ with the circuit represented by some parametrised unitary $U\left(\vec{\theta}\right)$, and some cost function $C\left(\vec{\theta}\right)$, Arrasmith et. al. [6] define the presence of barren plateaus as follows

Definition 2.19. A cost function $C\left(\vec{\theta}\right)$ exhibits a barren plateau if, for all $\theta_k \in \vec{\theta}$, the variance of the partial derivative of the cost vanishes exponentially with increasing qubit number, n, as

$$\mathrm{Var}_{\vec{\theta}}\left[\partial_k C\left(\vec{\theta}\right)\right] \leq F(n), \tag{2.20}$$

where $F(n) \in \mathcal{O}(\frac{1}{b^n})$

Arrasmith et. al. go on to demonstrate, via probability bounds and Chebyshev's inequality, that if the variance in the partial derivative of $C\left(\vec{\theta}\right)$ vanishes, then the probability that the partial dervivative is non-zero vanishes likewise.

Figure 2.13 showcases the difference between a cost landscape featuring a barren plateau, and a cost landscape that is trainable (no barren plateau). Note that a related phenomenon that often complements barren plateau landscapes is the *narrow gorge* at the optimal solution. A narrow gorge is the (very) low concentration of values for which

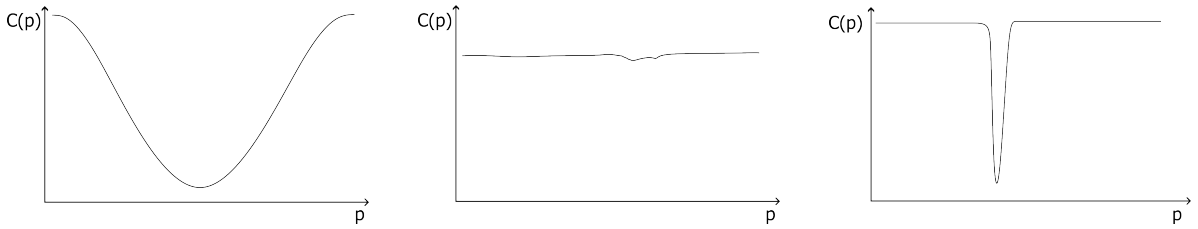

(a) (b) (c)

Figure 2.13: Simple, 2-D visual representation of cost function landscapes featuring different phenomena, with p as our parameter (space) that is being optimised over, and $C(p)$ the cost function we are optimising for. (a) A simple cost function landscape with no barren plateau, and easily findable minima. (b) A cost function landscape featuring a barren plateau with local minima and no narrow gorges. (c) A cost function landscape featuring a barren plateau with a single minimum found in a narrow gorge.

there is a significant gradient towards the minima, such that the probability of randomly sampling such a point from the underlying distribution is low. Due to the low number of samples 'within' a narrow gorge, even if an optimiser finds a value within one, it is likely to return to the barren landscape unless a very small step is used for optimisation; reducing trainability further. A more formal definition relating the probability of sampling a point within the narrow gorge to bounds in terms of dimensionality (qubit number) may be found in [6]. The same work also examines the permissibility of different cost function landscapes in parametrised quantum circuits, showcasing that only landscapes of the form in fig. 2.13 are possible.

Cerezo et. al. [19] established that local cost functions have only polynomially vanishing gradients, as opposed to exponentially vanishing gradients. Uvarov and Biamonte [111] explored this further, bounding the variance in terms of individual Pauli string coefficients (and their contribution the the overall variance). As a result, it is expected that the expressibility of an ansatz, in terms of local observables versus global observables, may help suppress the phenomenon of barren plateau, and help progress the field of VQAs.

The work presented in chapter 6 makes two contributions to the field: firstly, the mixed state compilation algorithm presented makes novel contributions to the field of VQAs, exploiting the result of the quantum low-rank approximation problems presented in [40] in order to provide compilation, compression, and principal component analysis all-in-one. Further, the work relating to local cost functions seeks to directly address the issues of barren plateaus.

Chapter 3

Experimental Methods

In chapters 4 and 5 we analyse the performance of two types of quantum-optical physically unclonable functions, *single-wavelength emission* and *hybrid-wavelength emission* quantum dot devices, respectively. The latter will henceforth be simply referred to as hybrid devices or hybrid PUFs, though they should not be conflated with the definition given to hybrid PUFs in [84].

These two types of PUFs are fabricated with an intent to exploit the uniqueness of optical PUFs for two different purposes. The single-wavelength emission devices are presented as a means for expanding PUF use and research to the realm of ubiquitous devices: due to the nature of quantum dot clustering in lacquer being capable of producing both micro-scale and macro-scale variations in emission patterns, it is hypothesised that such devices may be read with simple camera technology, such as that found in smartphones. The introduction of a second layer of quantum dot ink for the fabrication of the hybrid devices is expected to yield two different identities from a single challenge, with the identities being distinguished via the use of optical filters to extract information from only one ink's emission pattern. As such, these may not carry the potential for use with current-term ubiquitous devices, but instead may present novel ways to use PUF outputs for increased security in authentication, as well as cryptographic primitives. These potential use cases will be defined and further explored in chapter 5.

In this section, we detail the techniques used for the fabrication of both types of devices; the algorithms used to translate (classically captured) optical information from the devices in to secure authentication keys; and the processes by which we capture information from the devices. A simple overview of the interrogation process is shown in figure 3.1. Fabrication of devices was performed by Angelo Lamantia, with the processes and pipelines jointly determined by Lamantia and the author. As such, and due to the main focus being on measurement and interpretation of their optical properties, only a brief overview of fabrication methods is provided below.

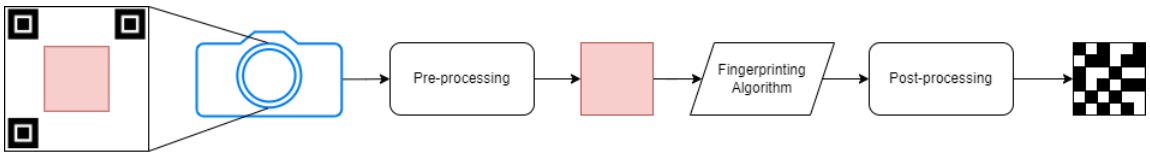

Figure 3.1: Flowchart displaying a simplefied view of the overall, generic processing pipeline for interrogating PUF devices in this work.

3.1 PUF Fabrication

For both single-wavelength and hybrid-wavelength devices, the uniqueness of the PUF device stems from the use of colloidal quantum dot-based inks. In each case, to fabricate the ink, CuInS/ZnS colloidal quantum dots are immersed in UVALUX UL-7M1000, a commercial UV lacquer, within a flat-bottomed vial at a concentration of 100mg/ml. With the use of a SciQuip Pro40 Digital Laboratory Mixer (equipped with a custom-designed 3D-printed tip), the lacquer-quantum dot combination is mixed at a constant speed of 750 rpm for 1 hour creating a dense and viscous mixture. Throughout, the flat-bottomed vial is immersed in an ice bath, negating the excess heat caused by the mixing process and thus preventing damage to the QDs themselves. The ink is then transferred to a flat surface via the use of a uniform hand-stamp, resulting in a 1cm^2 area of ink, which constitutes the PUF device.

3.1.1 Single-wavelength emission devices

The single-wavelength emission devices analysed throughout chapter 4 are primarily comprised of CuInS/ZnS colloidal quantum dots emitting at 650 nm, which were subsequently dispersed in a UV lacquer using a slow mixer at 750 rpm. The dots used have a specified core and shell thickness of 3.5 ± 0.5 nm, and 2.0 nm, respectively, with a concentration of 1.7 nmol/mg.

3.1.2 Multiple-wavelength emission devices

The devices analysed in chapter 5 were fabricated in a similar fashion, with the main difference being the addition of a second layer of quantum dots whose emission peak differs to the first set. Specifically, we use CuInS/ZnS colloidal quantum dots in both cases, first depositing a layer of quantum dots with an emission peak of 530 nm (CIS530), before depositing another layer of quantum dots emitting at 650 nm (CIS650). The CIS530 dots used have a specified core and shell thickness of 2.0 ± 0.5 nm, and 2.0 nm, respectively, with a concentration of 9.1 nmol/mg; whilst the $CIS650$ have the same specifications as described above. For 24 of the devices, the layer of CIS650 is deposited immediately after

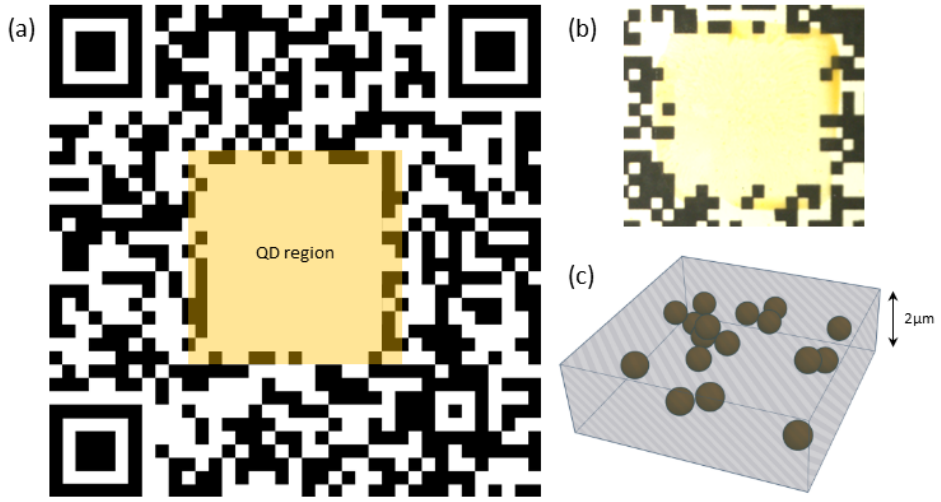

Figure 3.2: QD-PUF design. (a) Schematic of entire security device design, with the quantum dot-based ink PUF printed in the centre, occupying a 1×1 cm^2 QR-less region, and surrounded by a QR code which may be linked to the registration process. (b) Photograph taken of real device, cropped to include central area of QR in frame. The handstamped PUF is typically overprinted atop the device, ensuring that any edge artefacts (such as thick layers of ink seen in top right of QD area) are not fingerprinted. (c) Simple schematic of quantum dots suspended in UV-cured lacquer.

the layer of CIS530, resulting in an interface of two wet inks, before being left to dry. Another 24 devices were fabricated in which the layer of CIS530 is allowed to dry before the deposition of CIS650. The two groups of bi-wavelength emission devices, or, hybrid PUFs (H-PUFs) will henceforth be referred to by the monikers WH-PUF and DH-PUF respectively.

3.2 Lab-based data capture

In chapter 5, PUF devices comprising of two different types of quantum dots (with different wavelength emission peaks) are interrogated under laboratory conditions, making use of various optical filters in order to assess the separability of different fingerprints under the same challenge.

3.2.1 Shroud-based Capture

In order to image the entire optical PUF under well-controlled lighting conditions, a 'shroud' is used, comprising of a 3D-printed black box and lid, encasing a light source, CCD sensor, and a specially printed 3D stand for holding the PUF device. This setup is

pictured in Fig. 3.3. Each H-PUF device is imaged under two different conditions: using a combination of a long-pass 550 nm and short-pass 600 nm optical filters (to principally capture optical emission from the CIS530 dots); and using a combination of long-pass 650 nm and short-pass 700 nm optical filters (to principally capture optical emission from the CIS650 dots). A short-pass 400 nm short-pass optical filter is also placed in front of the light source (controlled throughout to provide constant illumination, with voltage and current maintained at 2.674 V and 0.163 A, via a T3PS3000 programmable DC power supply.), such that that no reflected incident light is captured in the images, whilst ensuring the incident light covers the absorption spectra of both dots used. We note that, given the emission spectra of both dots and potential 'leakage' from the filters, that it cannot be guaranteed that each case will only contain emission from the dot type being sought. However, such cross-over will be kept minimal (a point revisited in the presentation of results in chapter 5). The decision not to use fine, small bandgap filters was made as a compromise against the change in capture parameters that would be necessitated by such equipment: due to the scarcity of emitted light if the images were filtered with smaller (i.e., 10 nm) bandgaps, the exposure time used in capturing would have to be increased to the point where camera sensor noise would be increased, compromising the input image and thus the output fingerprint.

3.2.2 Microscope-based Capture

An optical microscope is also used, allowing for easier identification of areas dominated by either quantum dot for further investigation of how interaction between dots effects fingerprint commonalities, due to the clustering nature of the QD inks. Here, the general set up for data collection is outlined.

A *Zeiss Axio Lab A1* optical microscope is used, with a *Moticam S3* USB 3.0 camera attached directly to the microscope via the use of a camera mount adapter. All measurements are taken in brightfield imaging with epi-illumination, making use of an *Olympus MPlan* 10x lens. In the case of optical microscope measurements, captured light is again filtered twice: using a combination of short-pass and long-pass filters, as described above for the shroud case. However, the incident light is not filtered in this case, due to sensor noise that would be present in images taken with lengthy exposure times; as would be required given the low levels of light entering the aperture.

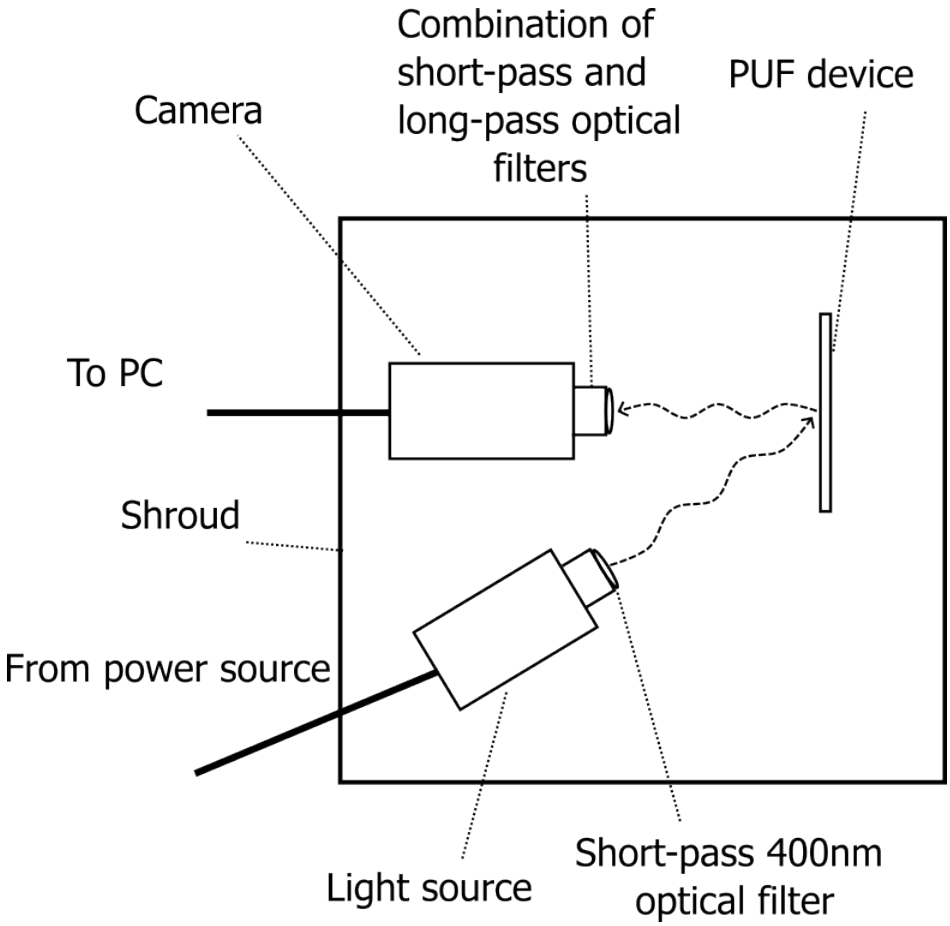

Figure 3.3: Simple schematic showcasing the shroud-based capture system. A small stand, upon which the PUF device is placed, is fixed to an optical bench, along with a light source, and camera. The wavy lines to and from the PUF device represent incident light and emitted/reflected light, respectively. A short-pass optical filter is secured in front of the light source, whilst individual pairings of short-pass and long-pass optical filters are secured in front of the camera, allowing for filtration of emission from each ink.

3.3 Smartphone-based capture

In order to assess the performance of our QD-Optical PUFs outside of laboratory conditions, we construct a framework for data capture using a smartphone; both for interrogation of the device (presenting a challenge), and capturing the unique response.

We make use of a specifically designed 3D-printed stand, to remove stability issues and retain a fine control of distance, L, between device and camera sensor, as well as the angle of incidence of challenging light, θ. A visual mock-up of the stand is shown in Fig. 3.5, with $L = 85mm$ and $\theta = 35°$. All images obtained via this set-up were captured on an iPhone X, with a constant ISO of 22, and exposure of 40 ms.

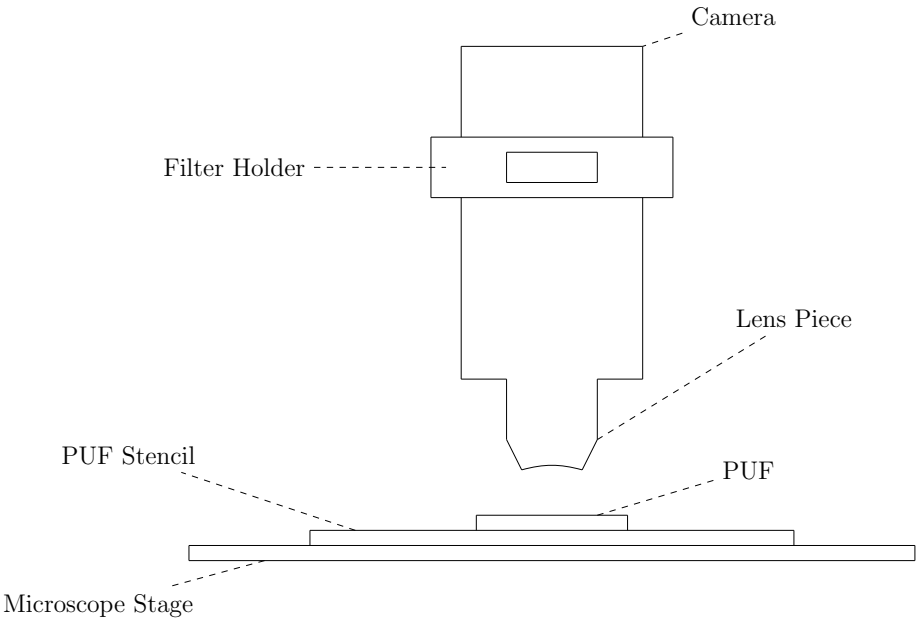

Figure 3.4: Simple schematic of the microscope-based capturing system. A PUF device is held still on the microscope stage via the use of a 3D printed stencil, ensuring no movement between captures, and that all captures are taken with the same orientation. A reflected light microscope is used, with epi-illumination brightfield imaging. Atop the microscope is fixed a filter holder, allowing for filtration of emitted and reflected light to capture different wavelength emission individually with use of a camera, screwed in on top of the filter holder.

3.4 Fingerprinting Algorithms

For the previously described PUFs to provide a means for authentication, it is not only necessary that they exhibit uniqueness when challenged; but we also need to be able to convert this uniqueness into some form of digital key for storage and verification. For this purpose, we deploy two different fingerprinting techniques; *Adapted High Boost*, and the novel *Reduced Modified Local Binary Patterns*. The former is a pre-existing algorithm developed specifically for optical PUFs, and serves to both benchmark the PUF devices, and the performance of the novel algorithm.

3.4.1 Adapted High Boost

Adapted High Boost (AHB) was first proposed as an alternative to the aforementioned Gabor Hash algorithm, specifically for the purpose of binarising information obtained from optical PUFs [98]. AHB takes as input a (single-channel) $a \times b$ pixel image, and outputs an $a \times b$ pixel bitmap, which encodes contrast information, as displayed in figure 3.6. AHB makes use of *convolution kernels*, a type of object regularly used in image

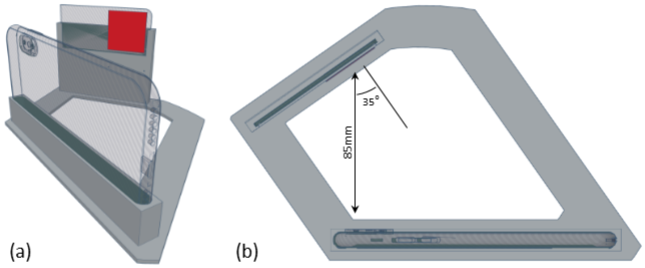

Figure 3.5: Visual schematic demonstrating experimental set up for capturing images of PUF devices using a smartphone. A PUF device is placed in a vertical stand, with the smartphone held in an opposing stand oriented to ensure a distance of 85 mm and an angle of 35 degrees from the PUF's normal plane for incident light.

processing, both for computer vision related tasks [98] and aesthetic reasons, with a host of photography post-processing toolboxes utilising them. Informally, a convolution kernel can be thought of as an $M \times N$ matrix "window", K of numeric elements, which, when "passed over" a section of an image, I, outputs a scalar value determined by some operation $\star$ between K and I. In practice, the kernel used in AHB (to be defined shortly) calculates the arithmetic mean of the intensity of all pixels within a neighbourhood of the center pixel, with the neighbourhood size determined by some choice of algorithmic radius, r. If the intensity of the center pixel (optionally, with some added, adjustable offset term) is smaller than this calculated mean, the pixel's corresponding output in the binary fingerprint is set to 1, otherwise it is set to 0.

Formally, AHB computes a given pixel's output using the convolution kernel

$$K = \begin{bmatrix} 1 & \cdots & 1 & \cdots & 1 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 1 & \cdots & -(n^2-1) & \cdots & 1 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 1 & \cdots & 1 & \cdots & 1 \end{bmatrix} \tag{3.1}$$

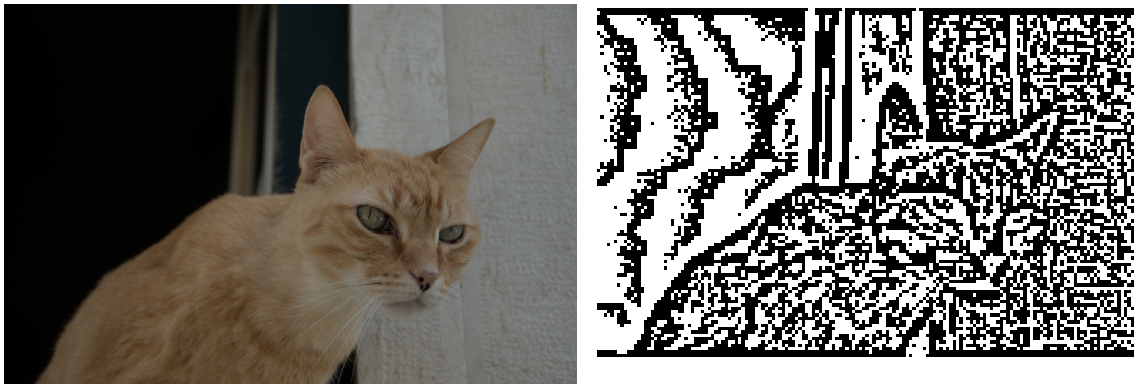

(a) (b)

Figure 3.6: (a) Image of a Portuguese cat. (b) Image shown in (a), converted to a single channel grey-scale image using MATLAB's native `rgb2gray` function and no resizing, and subsequently converted into a binary image via AHB, with a radius of 24.

and performing $K \star I$, where here $\star$ denotes element-wise multiplication followed by summing all elements, i.e.,

$$K \star I = \sum_{i}^{a} \sum_{j}^{b} K_{ij} \times I_{ij}, \tag{3.2}$$

and thresholding the output value, such that $K \star I < 0$ yields an output binary value of 1, and a binary value of 0 otherwise. For pixels near the boundary of the input image, from where the kernel may extend beyond the input boundaries, pixel values are interpolated via the standard 'replication' method.

The final fingerprint is then an $a \times b$ logical matrix, where each pixel's value is determined simply by whether or not it is brighter than the (average of) neighbouring pixels defined by r. As such, AHB can be seen to encode *local contrast information* from the original image.

3.4.2 Reduced Modified LBP

The second algorithm introduced is *Reduced Modified Local Binary Patterns*, or, *R-MLBP*. In contrast to AHB, R-MLBP principally encodes gradient information to output a binary image, as demonstrated in figure 3.7. The algorithm takes its name and fundamental ideas from Local Binary Patterns (LBP) [89], a texture description algorithm first formally described[1] in 1994, which has since seen great success when used for facial recognition [116], object recognition [100], and texture classification. The original LBP algorithm takes as its input a (typically one-channel, greyscale) $a \times b$ pixel image, and outputs

[1] In fact, LBP is a specific instance of the *Texture Spectrum* model [58].

(a) (b)

Figure 3.7: (a) Image of a Portuguese cat. (b) Image shown in (a), converted to a single channel grey-scale image using MATLAB's native `rgb2gray` function and no resizing, and subsequently converted into a binary image via R-MLBP, with a radius of 24.

an $a \times b$ *feature vector*, providing a characterization of textural information within the original picture. The algorithm functions (on a pixel-wise basis) in the following manner:

1. For a given pixel, store the values of it (α_0) and its eight neighbouring pixels (α_i for $i \in [0, 8]$ in a 3×3 matrix.

2. For each neighbouring pixel i, assign it the label '1' if $\alpha_i \geq \alpha_0$ or '0' otherwise.

3. Omitting the central pixel, concatenate the eight binary labels to form an $8-$bit binary number, α, known as a *texture unit*.

4. Convert α into decimal.

By storing the output decimal number generated for each pixel, an $a \times b$ feature vector is created; whose elements describe one of $2^8 = 256$ possible texture units. Each texture unit encodes local contrast information positionally. The original LBP algorithm as described above may be generalised beyond 3×3 coverage with 8 neighbouring pixels, instead defining the coverage by some arbitrary radius parameter, R, and choosing $N > 8$ 'neighbouring' pixels on the circumfrence of a circle centred at the centre pixel, with radius R. However, each calculated feature vector is prone to noise. Given that all information encoded in a feature vector is reliant upon the value of the centre pixel, a (very) noisy centre pixel may corrupt several output pixels away from the correct value, whilst noise in any neighbouring pixel may corrupt the value for the corresponding bit in the feature vector. Thus, to minimise noise, we make use of a more noise-resilient adaptation of LBP, *Modified Local Binary Patterns* (MLBP). MLBP, like the generalised LBP, takes two parameters: the algorithmic radius R and (an even) number of neighbour pixels,

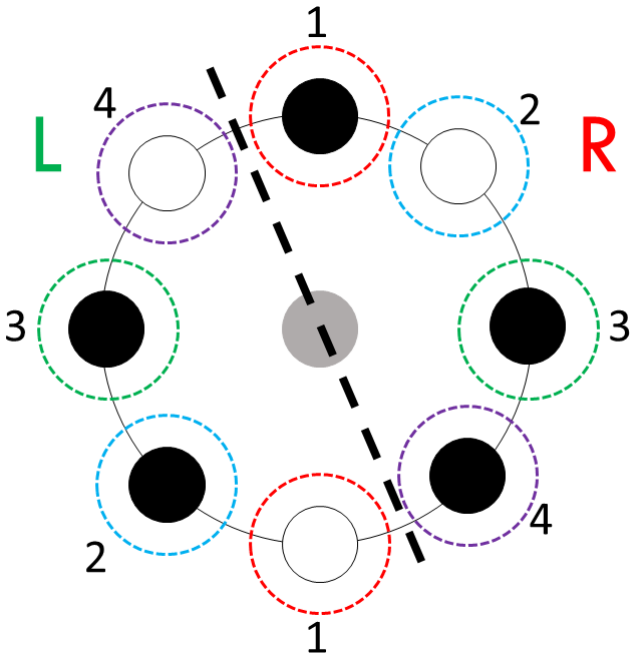

Figure 3.8: Simple representation of pixel-wise comparison points for RMLBP.

N. For a given pixel, the 'circle' of neighbouring pixels is split into a left half and right half, with the right half being defined as the first $N/2$ pixels going clockwise from the neighbouring pixel on the positive y-axis (upper pixel) from the centre. The algorithm then calculates a $\frac{N}{2}+1$ feature vector as follows:

1. Starting from the upper pixel and working clock-wise; compare diametrically opposed pixels. If the pixel value defined as being on the right half of the circle is greater than or equal to the value on the left, assign a bit-value of 1. Otherwise, assign 0. Concatenating these values provides the first $\frac{N}{2}$ bits.

2. Calculate the mean, μ, of all N neighbouring pixels.

3. Compare the centre value to the calculated mean μ. If it's greater than or equal to the mean, assign a value of 1. Otherwise, assign 0. This forms the $\frac{N}{2}+1$'th bit.

An example of this process for $N = 8$ is shown in figure 3.8. MLBP reduces the feature vector's dependency on the centre pixel in calculations (here, only contributing to 1 bit, as opposed to all bits as in LBP). Each neighbouring pixel only contributes to two feature vector bits, with the contribution to the final bit being minimal via averaging. By doing this, the algorithm shifts from encoding (positionally-informed) local contrast information, to instead encoding a mixture of local gradient information (from left to right, for the first $\frac{N}{2}$ bits) and (positionless) local contrast information (from the final bit). The result is a (compartively more) noise-robust feature vector than that of the

original algorithm, reducing the effect of both camera noise and inconsistencies in ambient lighting conditions during capture.

In order to yield an algorithm that maps each image pixel to a single binary value (as in AHB), we seek some binarisation, or *reduction* of MLBP's feature vector, motivating the use of *Reduced-Modified Local Binary Patterns*. R-MLBP adds one step to the above algorithm:

5. Take modal value of MLBP feature vector as final pixel output.

Which motivates the use of $\frac{N}{2} + 1$ values in the feature vector, and further reduces the effect of any pixel noise on a given (centre-) pixel's output.

Similarly to the implementation of AHB, towards the edges of the input image, a replication interpolation method may be used.

3.5 Processing Pipeline

Having detailed the algorithms used for fingerprint generation, as well as the systematic methods used to obtain the raw optical data for generating such fingerprints, here is detailed a (generalised) pipeline for the processing of information from the PUF devices. The framework follows the following steps:

1. Data capture — as outlined above.
2. Pre-processing — the processing of the raw captured image before input into AHB or R-MLBP.
3. Fingerprint generation — as outlined above.
4. Post-processing — the processing of the output of (either) fingerprinting algorithm, for i.e., noise reduction.
5. Output — the final output key, a 64×64 binary matrix.

A diagramatic overview is given in Fig. 3.1. As the earlier sections of this chapter have given great attention to the data capture and FP generation steps, here we provide information on pre- and post-processing.

3.5.1 Pre-Processing

Before passing an image to one of the two fingerprinting algorithms, we deploy a few pre-processing steps in an attempt to yield consistency in fingerprint generation and (to

some degree) to reduce noise. Some steps are tailored to specific algorithms, dot types, or capturing processes, whilst others are applied regardless. A general framework of the pre-processing is as follows:

1. Region of interest (ROI) cropping
2. Contrast boosting.
3. Channel selection.
4. Downsizing.

3.5.1.1 ROI cropping

All captured images initially include information surrounding the QD area itself (typically the QR code printed on the device). To ensure only the QD area (i.e., region of interest) is converted into the final output key, all images are cropped to square $c \times c$ images, containing only optical information from the QD emission itself. In the case of smartphone captures, due to the use of secure stands to hold the smartphone and the QD PUF, this process is streamlined by simply employing the same crop for all images taken. However, in the case of images captured underneath the optical microscope; due to the possibility of stage movement (as well as focal change) between uses, this necessitates that all intra-comparisons are only made between images captured in the same session (allowing one to ensure that the stage will not be moved between interrogations). Further, in such cases, regions of interest are determined and noted down, with specific cropping coordinates being used for each ROI.

3.5.1.2 Contrast boosting

Given the nature of AHB as a (local) contrast-encoding algorithm, simple contrast boosting is used to better highlight contrasting pixels and yield more reliable, unique fingerprints. To achieve this, the MATLAB *histeq* function is applied to all images before input to AHB.

3.5.1.3 Channel selection

For cases where a multi-channel camera sensor is used (producing RGB images), an important thing to consider is precisely what information to pass through the fingerprinting algorithms, which only accept single-channel images as inputs. Two factors influence this decision: the perceived colour of the quantum dot emission, and the specifics of the camera sensor itself. Naturally, for dots that emit red light upon excitation (i.e., the dots

used in this work whose emission peak lies at 650 nm), it would be expected that the blue channel of an RGB image would likely contain little to no valuable information from the dots emission, and would be more sensitive to noise. It might be expected that then the red channel should be used, however, in the case of images captured using a smartphone, it is standard practice in such camera development to have more sensors that capture green light, typically according to the Bayer array [10], due to the understanding of how human eyes perceive colour (linked with the fact that such cameras are often designed to produce the most 'aesthetically pleasing' result for human vision, as opposed to a true representation of the optical information of a given object). As such, for the results presented in chapter 4, the green channel is used for processing (note that, due to the (relatively) large FWHM of the emission spectrum of the CIS650 dots (see figure 5.2), the green channel will still contain a large amount of relevant information, even for dots perceived as red). For the images captured using a three-channel camera mounted to the optical microscope, a variety of channel selections are used in conjunction with the optical filters, for singling out individual dot emissions. For information captured between 550 - 600 nm (i.e., when using the long-pass 550 nm filter combined with the short-pass 600 nm filter), the blue channel is processed for fingerprinting, whilst for information captured between 650 - 700 nm, the red channel is used. For examining how well the fingerprints may be differentiated based solely on sensor channel filtration, information captured across the visible light spectrum (i.e., with no optical filtration used), information is processed from both the red and blue channels, as well as single-channel greyscale image extracted via MATLAB's native rgb2gray function.

3.5.1.4 Downsizing

Prior to fingerprinting, the resolution of the input image is reduced. This has two benefits for the processing of fingerprints: (i) It induces a speed-up of the algorithms (reducing the number of pixels (and hence number of iterations of the fingerprinting algorithm) processed); (ii) Downsizing of the input image typically reduces the amount of noise in the image, especially for single-pixel noise. Parameters for downsizing vary, depending on both the resolution of the original captured image, and the expected size of distinct features on the PUF surface. Parameters, as well as their choosing, will be defined where appropriate.

3.5.2 Post-Processing

The steps for post-processing are kept minimal, in an attempt to depend upon the directly captured optical information for fingerprint comparisons as much as possible. In fact, the

only post-processing step incorporated in the work of this thesis is for noise-reduction, and consistency in fingerprinting outputs; and that is resizing of the output. For both algorithms investigated, an $a \times b$ image yileds an $a \times b$ output. In order to reduce the effect of noise on the algorithm, all $a \times b$ outputs are resized to yield a 64×64 bit output, equivalent to a 4096 bit binary key, well above the standard thresholds for a secure public key. It is worth noting that, for real-world applications, there are a number of post-processing steps which are likely to be considered for use. For completeness and for consideration for future work focusing on applications of optical PUFs, a few are listed. They are not used here, as for this work it was considered important to understand the direct output of the fingerprint process as best as possible.

To attempt to lessen the effect of noise on output fingerprints, instead of only taking a single measurement of the CRP at a given time, multiple measurements may be taken and processed. Following the computation of several fingerprints associated with a given CRP probing, they may be combined to identiy anomalous pixels influenced by noise in the processing.

Additionally, for real-world security applications, a hash (or family of) algorithms is likely to be used to obfuscate the final output, thus making reverse-engineering of any given key harder, further complicating the process of gathering information in attempt to create a convincing fake device. However, the use of standard hashing algorithms alone is not suitable for the noisy readout of a PUF; due to the very nature of hashing, the Hamming distance between each output is likely to greatly increase, in comparison to the distance between each input fingerprint. Thus, a method is needed to obtain a consistent output key from all fingerprints obtained from a single device, despite the noise issue elemental to PUF readout. Numerous solutions for this exist, with fuzzy extractors [34, 29, 67] being a noteable example.

3.6 PUF Assessment

For the purposes of assessing PUF efficacy, as well as the comparison of different PUFs and fingerprinting algorithms, we desire a set of figures of merit to produce for each PUF instance. Throughout pre-existing literature, a number of different metrics have been proposed, but a single systematic set of metrics has not been adopted in a widespread fashion. Whilst this could be due to the (relatively) nascent nature of PUF research, we must also factor in that, with different challenge-response mechanisms characterising different PUF types, it is natural that different sources of (different types of) information may require varied approaches in assessment. In the case of PUFs whose output is naturally binarised, it may be sufficient to base the analysis only on the direct output

of the PUF itself. However, in the case of optical PUFs, the need for some (binarising) key extraction protocol naturally leads to the assessment being intrinsically tied to such a process. In this section is defined the system for evaluating our devices.

3.6.1 Assessment Model

For the analysis of PUF performance, we require a set of assumptions on any protocol that may use them. Here, we lay out the assumptions made during the research and writing of this thesis. Note that any real world implementation may involve a different set of assumptions, and some assumptions made here may be strengthened, or even relaxed, for such implementations.

1. *Verification process* — There exists some (black-box) process by which a device may be interrogated, and the process outputs ACCEPT (in the case that the device is real) or REJECT (in the case that the device is a fake).

2. *Fabrication method* — All (real) tags are fabricated via some prescribed method, which may be known to an attacker.

3. *Storage of keys* — It is assumed that all keys (in this case, the digital fingerprints) will be stored in a central system, for use within the verification process. It is further assumed that this information is kept private (although later, scenarios in which this latter assumption may be relaxed will be discussed).

4. *Fuzzy information error* — It is assumed that, due to the noise inherent to the signal capturing process, as well as uncontrolled/uncontrollable capturing conditions, different captures of the same PUF device will differ according to some error ϵ_{input}.

5. *Fuzzy fingerprint error* — It is assumed that, post-fingerprinting, the binary output key for a given PUF device should be equal to other output keys from different captures of the same device, up to some error factor ϵ_{FP}, dependent upon ϵ_{input}.

6. *Query model* — An attacker may query the verification process with fakes. In the case where the attacker may obtain a real device, they may also query the process with a real device. (The attacker will not gain access to the fingerprint associated with their device instance, only ACCEPT or REJECT.)

In the analysis presented in chapters 4 and 5, a full verification process is not considered, and is left for future work focusing on real world implementations. Instead, the focus is on the assessment of metrics that relate fingerprints (or a set thereof) to other fingerprints (or sets thereof). A full verification process would inherently need to

be combined with the query model and the implementation of key storage to ensure a cryptographically secure process, with safeguards against side-channel attacks as well as replay attacks. Thus, the discussion of side-channel attacks will be omitted from this work. However, replay attacks will be considered within our framework.

With regards to replay attacks, we consider three attack scenarios of an adversary, Malory, attempting to counterfeit a specific PUF device, X, with challenge-response (post-fingerprinting) output of K_X:

1. Malory has no information of a PUF's given input or output, and may manufacture a device, Y', using any material.

2. Malory has access to both a PUF's input image, and output fingerprint, and may manufacture a device, Y'' using quantum dots, mimicking the steps specified in 3.1.

3. Malory has access to both a PUF's input image, and output fingerprint, and may manufacture a device, Y''' using any material.

Scenario one simplifies to chance: Malory does not know the desired output she wishes to mimic, and as such, any device produced and interrogated for an output is essentially attempting to guess K_X. In a scenario where PUF outputs could be recorded repeatedly with perfect precision and no noise, the probability of passing is

$$\Pr\left[\text{Pass}\right] = \Pr\left[K_{Y'} = K_X\right] = 2^n,$$

where n is the total size of $K_{X,Y'}$. Given that the readout of optical PUFs is noisy, and authentication schemes require instead that a presented key need only be close to a registered key, in relation to some threshold parameter, t, opting to represent closeness of keys by the Hamming distance, we instead have

$$\Pr\left[\text{Pass}\right] = \Pr\left[K_{Y'} \sim K_X\right] = 2^{n(1-t)}.$$

In scenario two, Malory possesses information of both X and K_X, and as such may iterate their design approach, in order to produce a device, Y, that outputs a fingerprint $K_{Y''}$ which is sufficiently close to K_X. However, given that the production of quantum dot PUFs is considered to be a hard-to-control process, this scenario simplifies to attack scenario one: even though Malory has knowledge of what she must mimic in order to pass, as each quantum dot PUF device she produces has a random emission pattern beyond her control, the scenario can once again be considered as good as guesswork. In scenario three, however, the assumption that Malory will have little to no control over the optical pattern present on Y''' no longer holds. It should be noted that, given these devices

are interrogated via the use of a macroscale camera sensor, despite the origin of their uniqueness residing in the microscale, it may be possible to spoof the output response sufficiently enough to obtain a collision via the fingerprinting algorithm chosen; especially given the reliance on fingerprinting algorithms that are further distilling the microscale features into bits representing larger and larger areas of physical space. As such, Malory may expect she is able to produce a device, Y''', producing an output $K_{Y'''} \sim K_X$. However, in spite of this, such devices may still be used for secure authentication, by employing a secondary check that authenticates information about the material used, such as the non-linear emission response of quantum dots, as discussed in [44]. As it is understood to be impossible to replicate both the stochastic process of printing a device, as well as the clustering of quantum dots within the lacquer, it should be expected that, if a device can reliably be shown to exhibit these non-linear properties, then it's optical emission pattern could not have been highly engineered as a spoof attack. Thus, if such a material property check is implemented prior to authentication of $K_{Y'''}$, and devices which do not pass this check are rejected at this point, the issue of authenticating the output key simplifies to scenario 2, and can thus again be considered guesswork.

In this work, for each PUF CRP instance (detailed specifically for each case in chapters 4 and 5), the PUF will be interrogated under the associated challenge parameters multiple times, with the optical response being captured by use of camera sensors to generate a three-channel full colour images, as detailed earlier in this chapter. All images will be processed under R-MLBP and AHB, to create a fingerprint database. From this database, each CRP instance is assessed individually for both uniquness and reliability, by making use of the metrics defined in subsection 3.6.2. In order to do this, one image (and its associated fingerprint) are treated as equivalent to a registration entry; i.e., that entry will form the reference for comparisons made in the assessment of that CRP instance. Comparisons made between a reference entry, and entries from different interrogations of the same CRP instance will be dubbed *intra* comparisons, whilst those made between a reference entry and interrogations of a different CRP instance (and/or a different PUF instance) will be dubbed *inter* comparisons.

3.6.2 Performance Metrics

The following definition of metrics takes inspiration principally from the field of biometrics: given the noisy nature of the optical information, the problem of verifying that a given input matches a previously assessed given input must take into account this fuzzy element, similarly to the verification of biological fingerprints, irises, or indeed any biometric information. The nature of the task is to ensure that, given a previously registed PUF device, PUF_X, and fingerprint FP_X, any new captures of PUF_X should yield a small

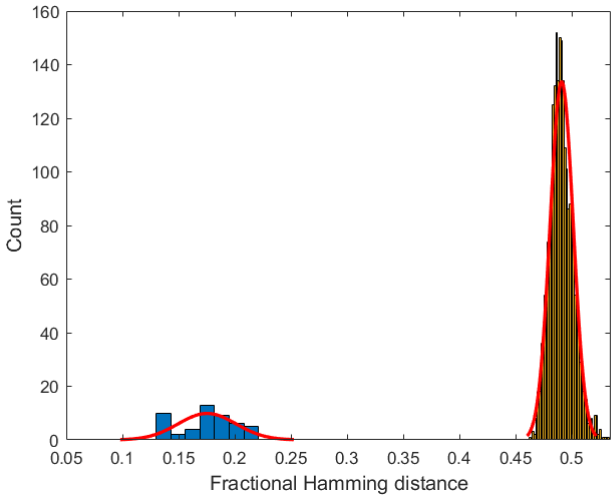

Figure 3.9: An example of the intra (left) and inter (right) Hamming distance distributions generated for a given PUF instance, by comparison against images of the same PUF instance (intra) and all other PUF instances (inter). The low x-values for the intra distribution, whose peak is just under 0.2, implies a reliable fingerprint extraction process over different images, whilst the peak centred around 0.5 for the inter distribution implies that the fingerprints retain the randomness of the PUFs' optical outputs.

value for ϵ_{input}, and the output of the fingerprinting process should emit a binary string $\text{FP}^*_X \equiv \text{FP}_X$ up to a small error ϵ_{FP}. Simultaneously, a capture of either a fake, or a different PUF instance, for now dubbed (w.l.o.g.) PUF_Y should yield a large value for ϵ_{input}, and in turn emit a subsequent fingerprint $\text{FP}_Y \neq \text{FP}_X$, determined by a large error ϵ_{FP}. In order to assess the similarity between different fingerprints (and therefore ϵ_{FP}, we make use of the *Fractional hamming distance*.

For a given PUF instance, by comparing the reference fingerprint to all other fingerprints from that instance generates the *intra-Hamming distance (intra-HD) distribution*. Likewise, by comparing the reference fingerprint to all fingerprints from other PUF instances generates the *inter-Hamming distance* (inter-HD) distribution. In an ideal case, elements of the intra-HD distribution would be equal to 0 (signalling two identical keys), whilst elements of the inter-HD distribution would be equal to 0.5 (signalling a difference on precisely half of all bits, implying seemingly randomness compared to each other). In practice, due to the fuzzy nature of the capturing process, this is not achievable, especially for the case of intra-HD elements. An example of both intra- and inter-HD distributions is shown in Fig. 3.9, with the separation between distributions clearly shown.

From the two HD distributions, both false acceptance rates (FARs) and false rejection rates (FRRs) for a given PUF instance can be calculated, via the cumulative distribution functions (CDFs) for each distribution. For the calculation of CDFs, a normal distribution is fitted to each HD set. Considering the intra-HD distribution for a given PUF CRP

instance to be a random variable X, and the inter-HD distribution being a random variable Y, we formally define the FAR and FRR as follows:

Definition 3.3. FAR

$$\begin{aligned} \mathrm{FAR}(X,Y) &= \Pr(Y \leq t_r) \\ &= \mathrm{CDF}(Y, t_r). \end{aligned} \tag{3.4}$$

Definition 3.5. FRR

$$\begin{aligned} \mathrm{FRR}(X,Y) &= \Pr(X > t_r) \\ &= 1 - \mathrm{CDF}(X, t_r). \end{aligned} \tag{3.6}$$

There are a variety of ways to consider how to choose the rejection threshold, t_r. Simply (but not necessarily practical for real-world implementations), t_r can be dynamically chosen independently for FARs and FRRs: for the former, $t_r = \max(X)$ whilst for the latter, $t_r = \min(Y)$. Cases where such choices are made are hereon dubbed *true FARs* and *true FRRs*. In such a case, the FAR and FRR respectively become

$$\begin{gathered} \mathrm{FAR}(X,Y) = \mathrm{CDF}(Y, \max(X)) \\ \text{and} \\ \mathrm{FRR}(X,Y) = 1 - \mathrm{CDF}(X, \min(Y)) \end{gathered} \tag{3.7}$$

For use in a wider authentication protocol, it is ideal to instead select a hard, fixed limit for t_r, that is jointly used in calculations of both FAR and FRR. Determining a suitable choice for this will be presented in the results section. As well as calculating FARs and FRRs, we choose a metric to more directly quantify the distance between the two HD distributions. Borrowing from the field of biometrics, we consider the *decidability* between two distributions, defined as follows:

Definition 3.8. Decidability

$$d' = \frac{|\mu_1 - \mu_2|}{\sqrt{\left((\sigma_1^2 + \sigma_2^2)/2\right)}} \tag{3.9}$$

The decidability is less arbitrary than FARs and FRRs, due to it being independent of any arbitrarily chosen parameter such as t_r. Further, it informs the decision on how to choose t_r: increasing t_r benefits the FRR, whilst worsening the FAR. As discussed in [26], a high decidability signals that we sacrifice less in terms of FAR whilst improving FRR, and vice versa.

Use of FARs, FRRs, and decidability help to test the separability of the HD distributions, as discussed, in order to assess how repeatable a fingerprinting process is, and the level of uniqueness retained in the final fingerprint from the initial optical output. However, none of these metrics provide insight on the information independently held within

an output fingerprint. As well as ensuring uniqueness and repeatability, we require a way to measure the inherent entropy in an individual fingerprint, as well as some checks on the content of the fingerprint to inhibit an adversary's ability to guess an output using some other publically available information about a PUF device, CRP, or fingerprinting process. One simple check is to consider the inherent *bias* towards a possible value of each fingerprint,

Definition 3.10. Bias. Given a binary $n \times m$-dimensional bitmap, α, we can calculate the bias of the bitmap as

$$\text{bias}(\alpha) = \frac{\sum_{i=1}^{m} \sum_{j=1}^{n} \alpha_{i,j}}{n \times m}, \tag{3.11}$$

i.e., the bias is equivalent to the proportion of ones in the bitmap.

Ideally, all fingerprints would yield a bias of 0.5, as large deviations from this imply an easier to guess output. However, in practice, perfectly unbiased fingerprints are unlikely to occur.

There are varying proposals of quantifying entropy of information in bitmap, whilst respecting spatial components of the image. Here, we opt to consider the degrees of freedom (DoF), initially defined as Effective Number of Independent Bits (ENIB) by Pappu [90], for a given fingerprint.

Definition 3.12. Degrees of Freedom. Given a binary bitmap (fingerprint) α, and a description of the random variable Y describing the inter-HD distribution calculated for α and a population of PUF CRP interrogations, $\mathcal{A}$ not containing α, one can calculate the degrees of freedom, N within α as

$$N = \frac{\mu\,(1-\mu)}{\sigma^2} \tag{3.13}$$

where μ and σ are the mean and standard deviation (respectively) of Y.

When assessing the above defined metrics, we (generally) seek the following criteria to be met:

- FAR under 10^{-6} and FRR under 10^{-5}, as a base for further use in an authentication protocol (specifically, for faith in the verification process).
- Decidability on par with biometric systems currently in wide-use, such as iris scanners, which yield a decidability of 14.1 in controlled conditions and 7.3 in field use.
- (For a key size of 4096 bits) A minimum DoF of 256 bits, to ensure that, even when deployed in an authentication protocol that further compresses the output,

the final key retains unique identifying bitstrings with a large information yield to avoid collisions (with or without hashing). A DoF of 256 bits is hereon considered a hard minimum, whilst aiming for a soft minimum of at least 512 bits to allow for some redundancy within the output.

The final list of metrics used in this work is:

- Intra- and inter-Hamming distances (HDs)
- False acceptance and false rejection rates (FARs/FRRs)
- Decidability
- Bias
- Degrees of Freedom

Given the (relative) nascency of the optical PUF field, a systematic method of assessing the performance of such devices is yet to be agreed upon. Due to that, the above metrics may not prove a definitive means of evaluating PUF performance. In [80], the authors propose a systematic method of evaluating PUFs, following a review of previous literature related to the assessment of PUF devices. In this, the following metrics are finally proposed:

- Uniformity
  - Defined to be the Hamming Weight of an output string, equivalent to bias. Ideally, a score of 50% would be achieved.
- Reliability
  - Equivalent to calculating the intra-Hamming distance distribution for a given tag. Ideally, devices are 100% reliable (equivalent to generating intra-Hammings of 0). Such scores are unachievable due to the noisy nature of the capturing process.
- Steadiness
  - Considered similar to reliability, but with a focus on capturing different conditions. Again, ideal target steadiness equates to intra-Hammings of 0. Small changes in the PUF's response in different lighting conditions are likely to propagate noise errors.
- Uniqueness

  - Equivalent to calculating the inter-Hamming distance distribution for a given tag. Ideally, this should correspond to inter-Hammings of 0.5, with scores less than this implying correlation, and scores above implying anti-correlation.

- Diffuseness
  - Considered similar to uniqueness, but instead considering differences between identifiers obtainable from a single device, hence being more applicable towards strong PUFs.

- Bit-aliasing
  - Defined as the Hamming weight of a particular bit across multiple devices.

- Probability of misidentification
  - Calculated via FARs and FRRs.

Comparing this list to the metrics used in this thesis, the majority are included. Of those that are not, two (steadiness and diffuseness) are due to their suitabilities, as stated in the report. Regarding bit-aliasing, as it focuses on recognising spatial irregularities in the outputs, it was deemed less relevant due to the inherent spatial uniqueness of the analysed PUF devices.

Further, decidability and DOF are used in this work, due to their suitability to weak PUFs and proposed use cases of this work. Decidability proves useful for assessing in-field use of the devices, allowing for confident matching of correct fingerprints given a potentially large number of devices. DOF provides a means of retaining confidence in the likelihood of a collision in the field being low, given a high number of devices in use. Both figures of merit allow for confidence in long term use of the devices, which is key in ensuring longevity in the use, and potential reuse, of a single weak PUF device.

Chapter 4

Quantum Dot Physically Unclonable Functions

The use of unique identifiers has been an essential cornerstone in the building of physical security devices, allowing manufacturers of goods, pharmaceuticals, and physical money to give a consumer confidence in the authenticity of their product. The battle against counterfeiting, however, has often taken the shape of a cat-and-mouse game, with manufacturers devising increasingly complex fabrication methods for devices, which only remain secure as long as the fabrication process is inaccessible to counterfeiters, typically due to financial and/or time constraints. Such a process (where, given the same resources, it is approximately as easy to create a convincing fake as it is to create a real, authentic device) is considered a *symmetric* manufacturing process. Due to the stochastic processes exploited by physical unclonable functions (PUFs) for their uniqueness, they can be seen as having an *assymetric* manufacturing process, i.e., even with full access to the equipment and knowledge required for device fabrication, it is (relatively) much harder to produce a convincing fake than it is to fabricate a real device; breaking the typical cat-and-mouse process that is typical of the counterfeit industry. However, the widespread adoption of PUFs has been limited due to the often complex methods, and expensive, inaccessible equipment, required to successfully interrogate a device and record its response.

In this chapter, the interrogation of optical PUFs utilising ubiquitous technology, namely a smartphone, is demonstrated in an attempt to surmount the roadblock of inaccessibility for PUFs. Quantum dot PUFs (QD-PUFs) consisting of CuInS/ZnS quantum dots with peak emission at 620 nm are assessed and shown to yield sufficient uniqueness to form a physical authentication token, with reliable interrogation and verification performed using only a smartphone's built-in torchlight (camera flash) and camera sensor. Firstly, this is demonstrated in ideal conditions, with the devices assessed in the presence

of no ambient lighting, ensuring the devices uniqueness can be reliably interpreted via the smartphone camera. Such conditions would also bear applications 'in-field', for a scenario where a product (e.g. a passport or identity card) utilising a QD-PUF may be assessed under dark conditions, such as a third-party (e.g. an immigration officer) placing the product into a trusted scanner. Working under the hypothesis that, without a strict control on ambient light, the flash from the smartphone is likely to be the dominating source of light over the device, the control on ambient lighting is removed. By interrogating the devices in a series of different lighting conditions and comparing outputs from each experiment, it is shown that the process can still be conducted reliably. This expands the potential use cases of a QD-PUF to the general case of allowing a consumer to verify the authenticity of a product in an unknown environment, without relying upon a third party for assessment.

In order to generate digital fingerprints from the captured optical emission, both reduced modified local binary patterns (R-MLBP) and adapted high boost (AHB) (as described in 3.4) are used on a dataset consisting of 48 PUF instances, each imaged 50 times. The resulting 2400 images are seperated into two classes:

- 48 images (one per PUF instance) which form the reference images; akin to the data used in the registration phase for a real-world implementation.
- 2352 images (49 per PUF instance) which form our comparison pool for intra- and inter-comparison populations.

Prior to any processing, crops of the captured quantum dot patch area are roughly 525×525 px, and are first re-sized to 175×175 px prior to fingerprinting. As detailed in chapter 3, output fingerprints are then resized down to 64×64 px. For these devices consisting of a 1×1 cm patch of quantum dot material, each pixel in an input to a fingerprinting algorithm corresponds to roughly 5.71×10^{-3} cm, whilst each pixel in a 64×64 output corresponds to roughly 1.6×10^{-2} cm

For both algorithms, each image is processed under 15 different choices of algorithmic radius between $r = 2$ and $r = 24$. For each choice of radius, one fingerprint is generated per captured image, resulting in a fingerprint population of 2400 per radius. When generating distributions for both intra- and inter-HDs, we only compare fingerprints amongst those obtained at the same choice of r.

Initially, the true FAR and true FRR are calculated (i.e., t_r is unfixed, and is adaptively changed according to the minima and maxima of the inter-HD and intra-HD distributions resp.). In the cases where the resultant calculation falls below MATLAB's precision threshold (roughly 10^{-308} for floating point double-type data objects), the output rate is rounded to zero.

4.0.1 Fingerprinting Outputs

For secure use in cryptographic protocols, the outputs of both fingerprinting algorithms (shown in figure 4.1) must satisfy certain requirements (as discussed in Ch. 3, to satisfy notions of randomness and uniqueness of the PUF itself, as well as reliability of the interrogation and fingerprinting processes.

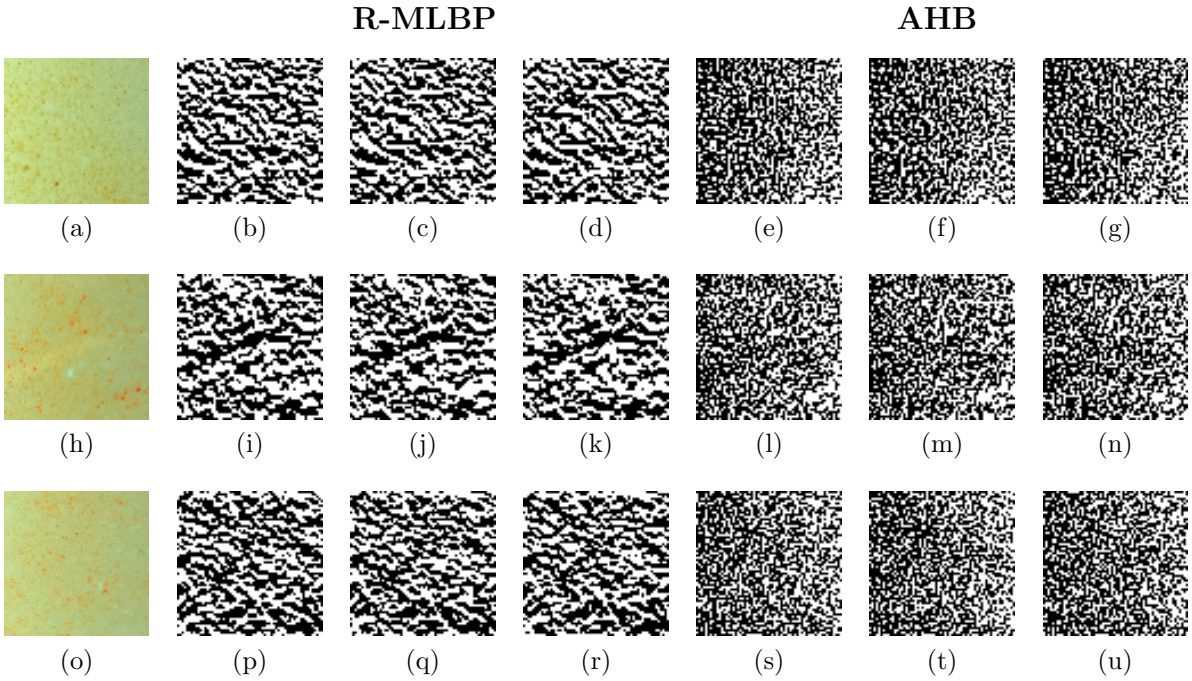

Figure 4.1: (a,h,o) Images of three different PUF devices. (b-d,i-k,p-r) R-MLBP outputs for the green channel of three different images of matching device; and (e-g,l-n,s-u) AHB outputs of the same three different images of matching device showcase both (i) the reliability of obtaining similar fingerprints from different captures of a given device; (ii) uniqueness of fingerprints of different devices. Each QD-region pictured in the left-hand column represents a $1 \times 1\,\mathrm{cm}^2$ region.

Figure 4.2 shows an example pair of Hamming distance distributions for a given PUF device under each algorithm, demonstrating both the uniqueness of a given device, and the reliability of the interrogation process.

As well as uniqueness and reliability, the outputs of the interrogation process must be considered random. Figure 4.3 showcases the evolution of the binary bias of fingerprints (on average) over changing radius. Both algorithms suffer from deteriorating bias as the algorithmic radius increases, suggesting that random variation within a device occurs at (relatively) small-scale levels. AHB achieves a worse optimal bias compared to R-MLBP. This poor performance in bias may be remedied (to some extent) by changes to the pre-processing, and/or tweaks to the algorithmic kernel.

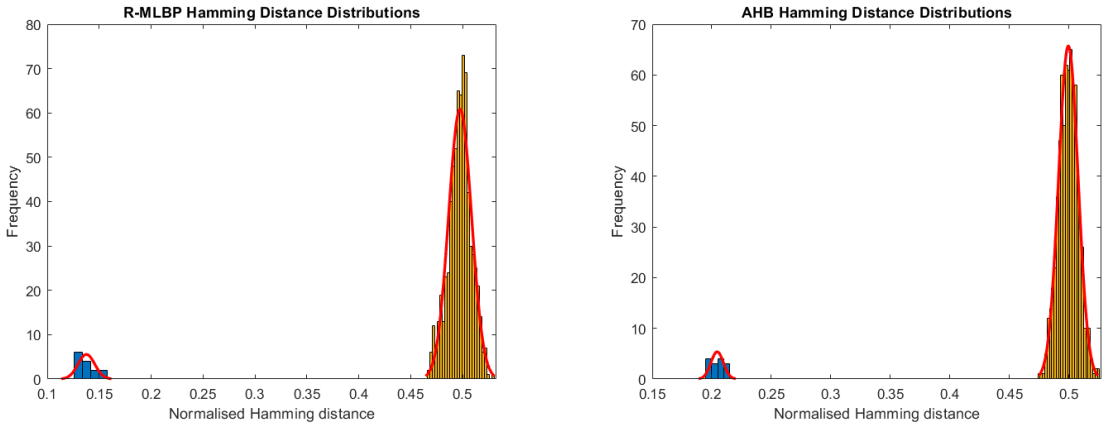

(a) (b)

Figure 4.2: Example intra-Hamming (left, blue) and inter-Hamming (right, orange) distributions for fingerprints produced across a range of QD PUFs under (a) R-MLBP and (b) AHB.

4.1 Shroud (Black-box) Conditions

4.1.1 Qualitative Comparison

When seeking to compare the two algorithms against each other, we find that both AHB and R-MLBP are capable of reliably producing unique fingerprints for all 48 QD-PUF instances. On average, AHB achieves an optimal absolute FAR of magnitude 10^{-179} whilst R-MLBP averages out at 10^{-170}. On 30 out of 48 of the tokens, AHB is deemed to yield lower absolute FARs than R-MLBP, initially suggesting that AHB is more capable of reliably producing fingerprints of the tokens, albeit only slightly. It is worth noting that, given the minuscule nature of the reported figures, such differences are trivial. Further, the reliance on the tails of distributions in calculating false rates should be treated with caution, as these tails may not be representative of experimental data. Despite AHB achieving lower optimal FAR values than R-MLBP, the latter yields a greater decidability (as shown in figure 4.4, suggesting it should be easier to find a fixed threshold for acceptance/rejection. R-MLBPs drop in FRR is therefore greatly beneficial when considering potential real-world implementations. In Fig. 4.5, we observe a much greater distance between produced FAR minima and maxima for R-MLBP, with FAR rapidly deteriorating as radius increases. This suggests that R-MLBP suffers from an inability to capture local information at high r, resulting in areas of uniformity within the output. As described in section 3.6.2, R-MLBP computes a final binary output dependent upon a fixed number, here, $N = 16$, of pixels on the circumfrence of a circle with radius r, regardless of the choice of r. Thus, with increasing r, R-MLBP calculates a deteriorating proportion of

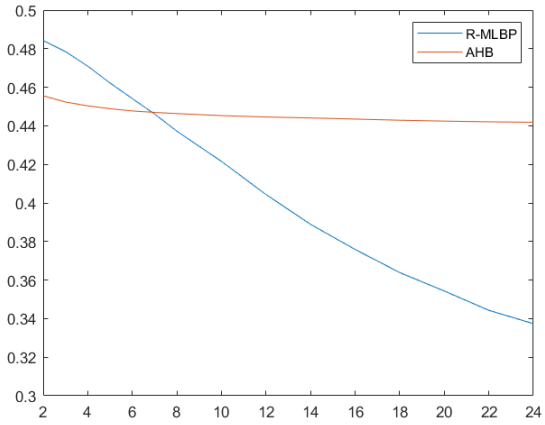

Figure 4.3: Plot of calculated average bias across entire range of fingerprints generated from PUF responses captured in all environments investigated in this work. 0.5 represents an ideal score for bias.

the total bits within the neighbourhood, compared to AHB, which always utilises all bits in the local neighbourhood. Consequently, R-MLBP's DoF performance also deteriorates as r increases. AHB retains locality and yields lower inter-hamming standard deviations, resulting in a higher entropy output across radii, as well as the aforementioned lower FAR performance. However, R-MLBP benefits from a lower intrahamming distance and standard deviation, resulting in much lower FRR values and greater decidability. Examining the choice of radius at which absolute FAR is optimized (i.e., r corresponding to a global FAR minima), we find that R-MLBP averages an optimal radius of 2.6, whilst AHB averages at 9.7, significantly higher (a trend further highlighted in Fig., where we see that R-MLBP's absolute FAR tends to increase with radius, whilst AHB's decreases). Noting that algorithm performance at higher choices of r involves the amalgamation of information from a greater pixel neighbourhood, we hypothesize that AHB is less-well equipped at reliably extracting the fine-grained uniqueness from QDPUFs, and instead excels at extracting coarse-grained uniqueness. Given the potential of manufacturing thousands, if not millions, of PUF devices for use in commercial applications, this difference in graining by the algorithms may prove pivotal in wide-scale robustness: PUF devices which prove uniquely different within small-pixel neighbourhoods may yield coarsegrained fingerprints that are more likely to form a collision under AHB (due to large(r) segments of repeated bits within the fingerprint), when compared to the excelling fine-grained fingerprints of R-MLBP. However, due to correlations inherent within outputs, it may be useful to couple fingerprint generation with some form of entropy addition process. Additionally, the standard deviation of optimal r is 0.84 for R-MLBP, and 5.72 for AHB. In a commercial application, this provides greater confidence in being able to choose a single radius to as-

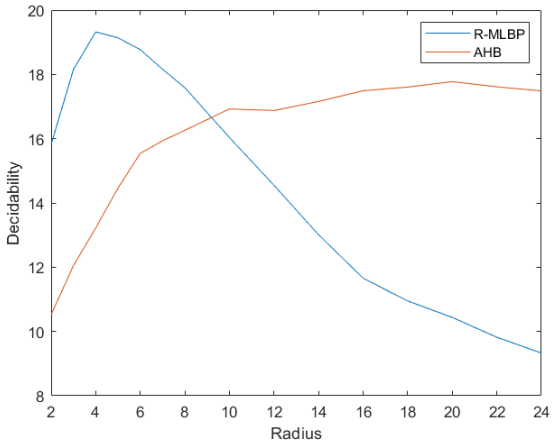

Figure 4.4: Decidability scores averaged over all 48 PUF devices investigated, shown over varying radii for R-MLBP and AHB. Whilst R-MLBP outperforms AHB, both algorithms scores are comparable to those reported for modern iris recognition technologies.

sess all produced QD-PUFs under R-MLBP than under AHB. Further, fingerprinting at higher radii is more computationally intensive under AHB, with the number of necessary calculations scaling up quadratically with r, which, considering implementation times for user experience in real world use, suggests seeking a low working radius.

4.2 Ambient Lighting

The results presented thus far serve as a proof of concept for the idea of measuring optical PUFs with ubiquitous devices. However, for widespread adoption, the process of interrogation and recording of the QD-PUF challenge-response pair (CRP) must be able to withstand a plethora of environmental factors, such as varying lighting conditions, movement (due to the human element of holding the smartphone and PUF device) and perspective correction. Towards this aim, this section presents an analysis of performance in different lighting conditions.

All 48 devices are re-interrogated in three different ambient lighting conditions: low-light (0.006 kLux) in an optics lab, and two medium-light (0.220 kLux and 0.326 kLux) in an office and an optics lab respectively. Initially, results are presented to show the reliability of obtaining a unique identifier in each individual setting, before comparing identifiers obtained from all settings against the earlier analysed identities from the shroud captures. The latter analysis is used as a model of a real-world implementation, in which an identifier obtained in any setting must be successfully authenticated against the identity obtained at registration (in this case, blackbox shroud conditions).

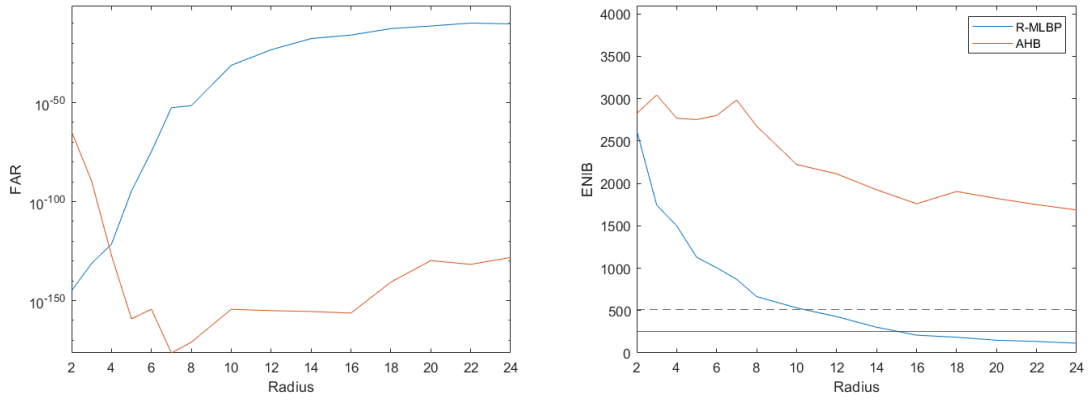

(a) (b)

Figure 4.5: A comparison of (a) false acceptance rate (FAR) and (b) expected number of independent bits for a single-dot emission PUF device, under both R-MLBP (blue line) and AHB (orange line). The dashed line in panel (b) represents the minimum allowed score for ENIB.

4.2.1 Single session comparisons

Initial assessment of the PUF devices in ambient lighting conditions is done on a per-session basis: the analysis previously done within blackbox shroud conditions (but utilising only the inner stand) is repeated in each of our chosen ambient lighting conditions. The aim of this analysis is two-fold: it allows for a proof that reliable fingerprints are capturable outside of shroud conditions, and it allows for an assessment on the impact of ambient light on intra-hamming distances, providing data on noise-resistance of each algorithm, as well as data for planning mitigations in order to improve performance.

For each environment, a light-meter was used to capture the level of local ambient light before commencing the PUF interrogation process. Two capturing environments were located in the same optics lab in which the shroud was used, but without the shroud, and varying ambient light via the combination of different levels of indoor lighting; resulting in a scene at 0.06 kLux and one at 0.326 kLux. The final environment was in an office with a combination of natural light and ambient light resulting in a local light-level recording of 0.220 kLux. The results are presented in order of increasing ambient light levels.

4.2.1.1 0.06 kLux

Outside of shroud conditions, we again find that both AHM and R-MLBP are capable of reliably and repeatedly producing unique fingerprints for all 48 QD-PUF instances. However, surprisingly, the minima of the (non-fixed threshold) FARs tend to be lower in ambient lighting conditions, as seen in Fig. 4.6: For AHB, the average global minimum of

the FAR curves across all PUF devices is of magnitude 10^{-284}, whilst FARs for R-MLBP average out at a magnitude of $10^{-}260$. Initial expectations were that FAR performance would worsen with the introduction of ambient light. However, in low levels of light, it is hypothesized that the improvement in performance will be related to the automated capturing process of the camera; with the camera's focus being maintained correctly in the presence of ambient light.

The average radii at which the global minima were found under AHB is $r = 5.3$, with a standard deviation of 3.17, whilst for R-MLBP, the average was found to be $r = 2.1$, with a standard deviation of 0.31

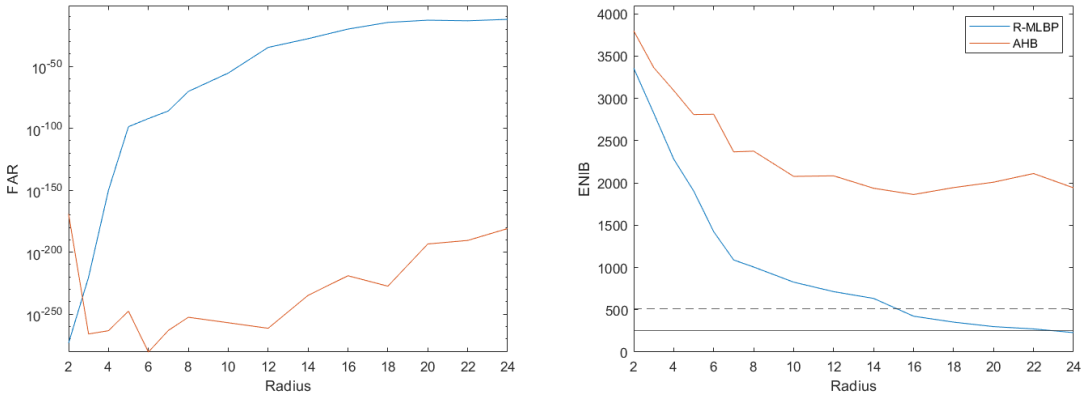

(a) (b)

Figure 4.6: A comparison of (a) FAR and (b) ENIB for a single dot emission PUF device. Both sets of results are calculated based on inter-Hamming distance distributions generated from fingerprints of all PUF devices examined in an open shroud, with local ambient light level readings of 0.06 kLux. The dashed line in panel (b) represents the minimum allowed score for ENIB.

4.2.1.2 0.326 kLux

Increasing ambient light levels further, the devices were re-interrogated in the same laboratory. Once again, the reliable production of unique identities is proven to be possible with both examined algorithms. The results directly follow those presented above for the 0.06 kLux case, with AHB again producing lower (non-fixed threshold) FARs than R-MLBP on 30 out of 48 PUF devices. The average of the global minima of the FAR curve for AHB is of magnitude 10^{-287}, and for R-MLBP, the average is 10^{-262}, showing similar improvments over the 0.06 kLux case. For AHB, these minima were found at an average radius of $r = 4.1$ with a standard deviation of 2.09, whilst for R-MLBP the equivalent figures calculated were $r = 2.1$, with a standard deviation of 0.33.

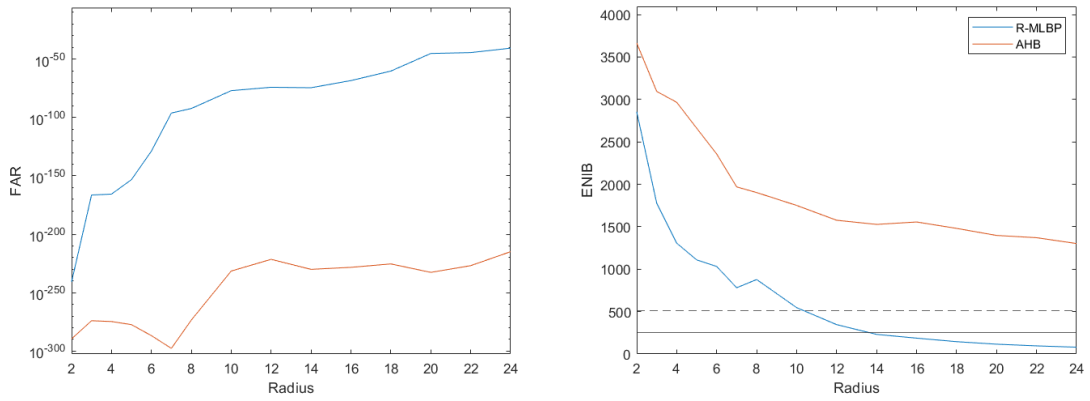

(a) (b)

Figure 4.7: A comparison of (a) FAR and (b) ENIB for a single dot emission PUF device. Both sets of results are calculated based on inter-Hamming distance distributions generated from fingerprints of all PUF devices examined in an open shroud, with local ambient light level readings of 0.326 kLux. The dashed line in panel (b) represents the minimum allowed score for ENIB.

4.2.1.3 0.220 kLux

Additionally, the devices were also interrogated outside of an optics laboratory, allowing for the influence (indirect) natural light on the capturing scene. From readings taken before and after the round of interrogation, the average level of ambient light throughout the session was 0.22 kLux. This shift away from well-controlled light introduces the first big change in the trend of results so far, with AHB continuing to benefit from the increased ambient light levels, whilst R-MLBP's performance decreased: the average optimal FAR for AHB (R-MLBP) in these conditions was of magnitude 10^{-285} (10^{-251}). These minima were obtained, on average, at a radius of $r = 4.6$ ($r = 2.1$) for AHB (R-MLBP), with a standard deviation of 2.2 (0.31).

4.2.1.4 Assessment

In each environment that the PUF devices were interrogated in, both fingerprinting algorithms were able to reliably fingerprint the captured image and obtain unique identifiers. In general, the results follow the trends established in the earlier, shroud-based work: AHB achieves lower optimal FARs than R-MLBP, whilst R-MLBP is more capable of separating the intra- and inter-Hamming distance distributions (as shown by the decidability scores shown in figure 4.10). Figure 4.9 showcases the similarity in trends for both algorithm's performance in these environments, with intra-(inter)Hamming distance distributions consistently low (~ 0.5) for each environment.

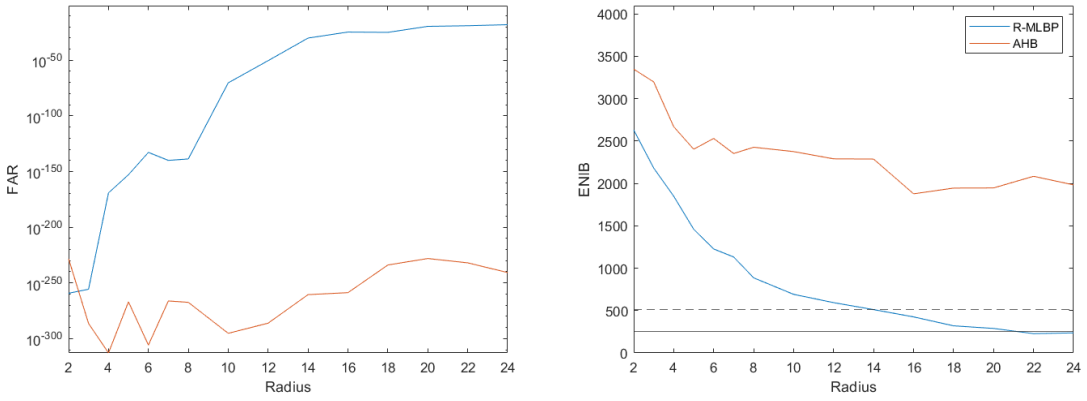

(a) (b)

Figure 4.8: A comparison of (a) FAR and (b) ENIB for a single dot emission PUF device interrogated under ambient light provided by internal office lighting and indirect external sunlight through a nearby window, with an ambient light reading of 0.220 kLux. A comparison of (a) FAR and (b) ENIB for a single dot emission PUF device. Both sets of results are calculated based on inter-Hamming distance distributions generated from fingerprints of all PUF devices examined in an office environment including both internal and external lighting. Local ambient light level readings taken before and after measuring averaged out at 0.220 kLux. The dashed line in panel (b) represents the minimum allowed score for ENIB.

Examining the radii at which optimal FARs are achieved across different lighting conditions, it is worth noting that R-MLBP's optimal performance was consistently focused around $r = 2$, with low levels of variation, in comparison to AHB, whose optimal performance is more variable across different conditions, focusing around $r \approx [3, 5]$. This further boosts the prospect of R-MLBP for a real-world implementation, due the increased assurance of good performance at one chosen fixed algorithmic radius.

4.2.2 Comparisons against shroud reference

The results shown in the previous subsection demonstrate the PUF device's ability to be reliably interrogated in different lighting conditions. However, for real world use, fingerprints generated from user-captured images of a device would be unlikely to be compared to those generated in similar lighting conditions, instead being compared to a reference fingerprint generated in controlled conditions at registration. Thus, it must be ensured that fingerprints captured in each lighting condition are adequately similar to those generated at some registration point. In order to test this, the earlier captured in-shroud images are treated as a reference point generated during some registration period, and these fingerprints are compared to those generated in different lighting conditions.

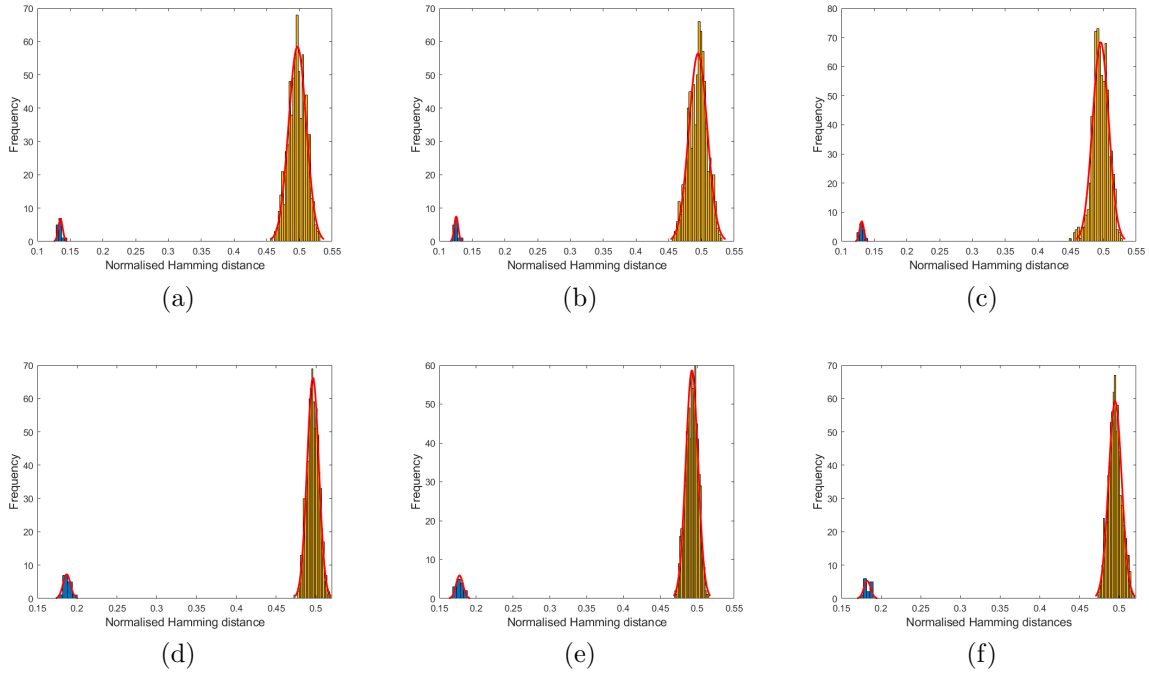

Figure 4.9: (a)-(c) Intra-Hamming (left,blue) and inter-Hamming (right, orange) distributions generated via R-MLBP from images captured at 0.006 kLux, 0.326 kLux, and 0.220 kLux respectively. (d)-(e) Intra-Hamming (left,blue) and inter-Hamming (right, orange) distributions generated via AHB from images captured at 0.006 kLux, 0.326 kLux, and 0.220 kLux respectively.

In this case, intra-Hamming comparisons are made between the reference image, and all images (of the same device) captured from each environment, resulting in 45 data points for the distribution for each device. Inter-Hamming comparisons are made between the reference image, and all images of other devices from each environment, yielding 2115 data points per device.

Examining decidability metrics (shown in figure 4.10) calculated across the data population, R-MLBP outperforms AHB, suggesting that a fixed threshold for separation of of fingerprints from the wrong (or, fake) devices is easier to find for R-MLBP than for AHB. In earlier sections of this chapter, "adaptive" false acceptance and rejection rates were shown, acting as a measure of distance between each Hamming distribution, and the nearest extrema of the other. Whilst these are of use for understanding the likelihood of an element of one distribution belonging to the other, such rates are unlikely to bear usefulness in real-world applications, where a fixed threshold for rejection or acceptance of a given fingeprint is required. To this end, for the model of an authentication protocol, fixed FARs and FRRs are calculated for the Hamming distributions; i.e., the probability of an inter-(intra-)Hamming distribution element falling below (above) some fixed threshold, t, is calculated and presented as the $t-$thresholded FAR (FRR). Results are

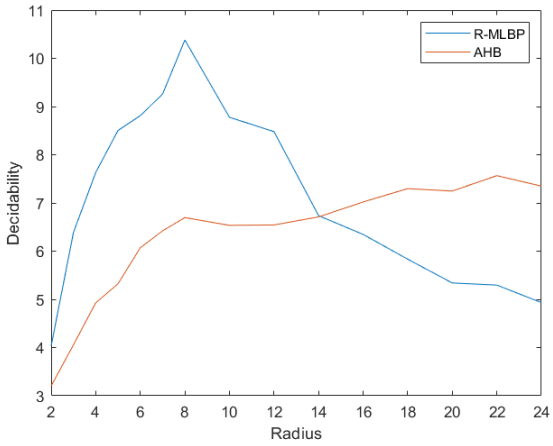

Figure 4.10: Example of the average decidability calculated for a PUF device, over varying radius for fingerprints generated via R-MLBP (blue) and AHB (orange). Hamming distance distributions used for the calculation of decidability were obtained from fingerprints of captures taken in all lighting conditions described.

shown in figure 4.11 for a selection of thresholds $t = 2.5, 3, 3.5, 4, 4.5$. As suggested by the calculated decidability scores, R-MLBP is better capable of achieving strong scores for both FAR and FRR at a fixed threshold for a single radius choice for both rates. Whilst AHB is again capable of producing lower FARs overall, for fixed thresholds it achieves much worse FRRs than R-MLBP throughout the choices of radius.

4.3 Concluding Remarks

Optical PUFs made using collodal CuInS/ZnS core/shell colloidal quantum dot-based inks are proposed as strong candidates for use in ubiquitous PUF designs. It is shown, through the low intra-Hamming distance distributions with mean less than 0.2, that they can be reliably interrogated in a variety of lighting conditions using a ubiquitous device, namely a smartphone, in order to both excite the quantum dots optically, and capture their emission pattern response. Further, two algorithms are shown to be able to reliably process captures into unique (in the case of different PUFs) fingerprints for further processing in cryptographic communications, highlighted by the inter-Hamming distributions produced, all with a mean around 0.5. Of the two, R-MLBP produces more consistent results across different capturing environments, suggesting it to be more suitable for use. It is found that a matching threshold of $t = 0.35$ would allow for easily differentiable fingerprints, with expected percentage of false acceptances (rejections) around 10^{-20} (10^{-5}) for a radius choice of 8. However, the reported rates should be taken with some caution. Whilst it is standard within PUF literature to used Gaussian fits to approximate the

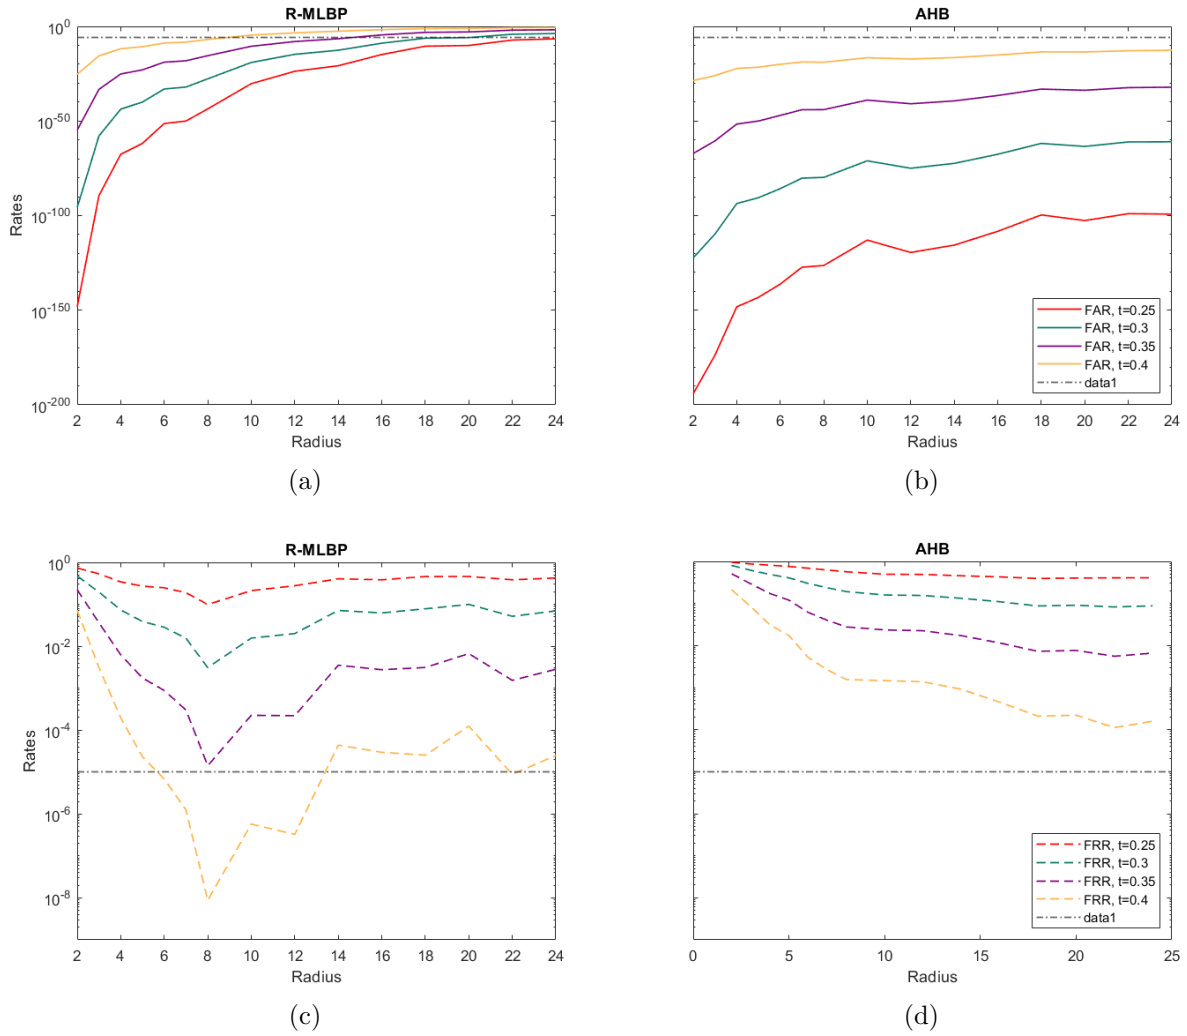

Figure 4.11: 4-panel figure of thresholded FARs and FRRs, for $t = 0.25$, 0.3, 0.35, and 0.4. (a) FARs for R-MLBP. (b) FARs for AHB. (c) FRRs for R-MLBP. (d) FRRs for AHB. The targets defined in chapter 3 (10^{-6} for FAR and 10^{-5} for FRR) are annotated via the dot-dashed lines.

Hamming distributions, this is often applied to electronic PUFs, for which the output bit string is expected to be independent and identically distributed (i.i.d.). However, due to the nature of optical PUFs, bits may not conform to the i.i.d. assumption, due to the detection of features of varying sizes by fingerprinting algorithms, as well as the fact that one bits value is determined by the input of pixels around it. As such, future work should investigate the use of other fits, and for any large-scale experiments, work can be done to compare these fits to a large pool of experimental data. Due to the lower reported intra-Hamming scores, it is believed that this strong performance across conditions is due to its inherent noise resilience. AHB particularly suffers from an inability to distill unbiased fingerprints from captured responses, further limiting its suitability. To this end, changes

to the post-processing in order to introduce some debiasing may benefit both algorithms, but particularly AHB.

Whilst the sample size presented in this preliminary study allows for reported ENIBs to be comparable to those presented in the original PUF work by Pappu, they fall orders of magnitude short of what would be necessary to replicate a real-world use case. And given the noisy nature of PUFs, the devices can certainly not be expected to fill the entire key space of 2^{4096}, and may fall short of filling the key space suggested by reported ENIBs. However, the chosen key size allows for a great deal of redundancy within the output, in order to mitigate for this, and whilst future studies with greater sample sizes may report a drop in ENIB, the results presented here remain positive for future development.

As mentioned, a limitation of this study is the use of the custom-printed stand, which fixes both the distance and angle between the capturing device, and PUF device, for smartphone-based data capture. For real-world use, however, both distance and angle of capture can be controlled (to some degree) via the design of the app used, utilising both textual and visual prompts. Further, the use of standard perspective correction algorithms for image processing may be used to reduce any impact on capturing angle that may be introduced by handling of the devices throughout the interrogation process. This is discussed further, with examples shown, in Appendix A.

Additionally, for long term use, and for robustness in different conditions (both in terms of consumer use, and packaging and logistics), any body wishing to use quantum dot based PUFs would need faith in their resistance to ageing, as well as environmental conditions such as heat. With regards to ageing, different use cases may have different requirements, for example, due to varying shelf-lives on products. The work in [78] examines the effects of both temperature and ageing on QD-PUFs, and the stability of their fingerprints under R-MLBP. Further, recommendations are made on ways to increase stability and arrest photodecay over time, allowing for further faith on products with long shelf-lives.

Chapter 5

Hybrid QD PUF

In this chapter, a novel form of quantum dot physically unclonable function (QD-PUF) making use of a second layer of quantum dots (with a peak emission different to that of the first layer) is proposed and assessed in order to yield two discernible fingerprints from the same optical challenge. Interrogation of the devices is performed in laboratory conditions, initially utilising a shroud in a similar vein to the work conducted in chapter 4 to display a proof of concept, yielding unique fingerprints from different PUF devices, and assessing the similarity from fingerprints formed from the optical emission of the different QDs on a single device. In order to isolate emission contributed from each ink independently, a series of short-pass and long-pass filters are placed in front of the camera, to focus the captured emission on a bandgap around the target dot's peak emission wavelength. A short-pass filter is placed in front of the incident light source in order to ensure no reflected light is captured from the exciting of the dots.

Due to a variety of factors, such as overlap of broad emission spectra of the quantum dots; overlap between the emission spectra of one, and the absorption spectra of the other; and spatial relations between the two inks arising from fabrication (giving rise to the potential for carrier transfers between nearby clusters, with interactions potentially varying further with excitation wavelength), it's expected that some level of correlation between fingerprints of different emission profiles from a single device may exist. Assessments of similarity show this to indeed be the case: two fingerprints obtainable from a single device are different, but not as unique as fingerprints arising from different devices. To investigate this phenomenon, devices are then assessed under an optical microscope, observing how spatial proximity of large clusters of QDs contributes to shared elements of output fingerprints.

This chapter is organised as follows: First, the notion of PUF strength is revisited, with a slight reformulation being introduced to capture the behaviour of these hybrid PUFs. Then follows a discussion of the devices themselves; their use as an extended

source of entropy in comparison to single-emission devices; and how expected correlations may contribute to increased security and use-cases. Results are subsequently shown, with a discussion following.

5.1 One-to-Many PUF

Thus far, the use of quantum dots in authentication devices has been considered in the scenario of a traditional weak PUF: we have only considered a single challenge-response pair (CRP) mechanism (with the aims of compatibility with ubiquitous technology), and thus, a trivially small challenge-response space. The typical view of a strong PUF involves the use of multiple challenges to generate multiple (unique) responses. This can be pictured as being akin to the notion of an injective, or one-to-one, function, mapping distinct elements of its domain (challenges) to distinct elements of its codomain (responses). Here, we propose a different formulation of a PUF, instead akin to a one-to-many map. Again, analysis will focus on the presentation of one challenge, but with the aim of capturing multiple, distinguishable, responses. To construct such a PUF, an ink

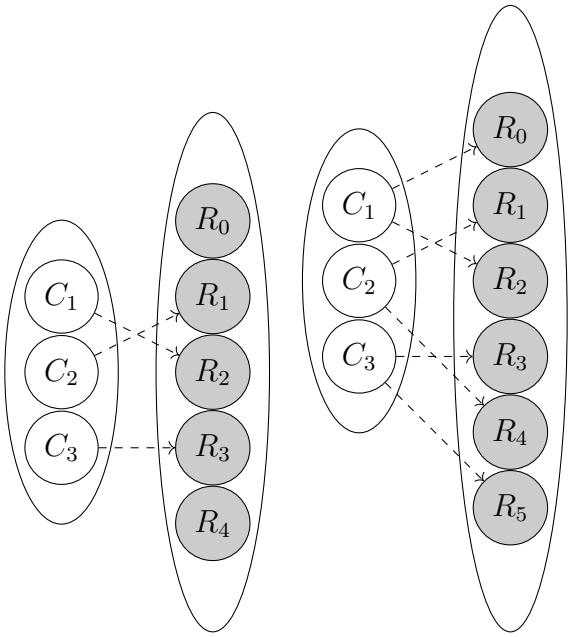

Figure 5.1: Diagramatic representations of (left) typical multi-CRP PUF (one-to-one) and (right) proposal of one-to-many multi-CRP PUF. Note that, in the work presented in this chapter, the different responses (e.g., C_1 mapping to R_0 and R_2) are distinguished via optical filters.

comprising of two different sizes of quantum dot molecules is used, with the responses being distinguished by only capturing light focused around each dot size's respective emission spectrum peak. We make use of two sets of CuInS/ZnS core/shell quantum dots,

with peak emission at 530 and 650 nm, with the PUF fabrication process detailed in Ch. 3. Similarly to the single-emission devices discussed in Ch. 4, the devices can be interrogated with a simple light device. By filtering the light that passes through to the capturing devices, it is expected that two unique fingerprints are obtainable; one from the emission of CIS530 quantum dots, and one from the emission of CIS650 quantum dots. Atop of this, a third fingerprint is also obtainable by choosing not to filter the captured light, and instead fingerprinting the emission of both CIS530 and CIS650 quantum dots. To some degree, this expansion of the CRP space may seen as equivalent to the segmentation of a PUF response, in which, given a large response, it can be segmented into multiple outputs and, thus, treated as multiple CRPs, similarly to [98]. However, one-to-many PUFs offer a potential advantage by allowing segmentation to occur for the same spatial points of a device. This has direct implications on discussion of PUF strength, where CRP concentration in comparison to device size is used as a measure; and may further help bridge the gap between PUFs and ubiquitous technology, allowing the devices to take up less real estate on products. An additional discussion on the potential further expansion of the CRP-space is provided at the end of this chapter.

5.2 Correlated Hybrid Extended Entropy Source

As is the case for the single-dot emission devices discussed in chapter 4, the uniqueness of emission pattern for a deposited layer of the quantum dot ink stems from the random and uncontrollable clustering of quantum dots in laquer, with an additional contribution coming from the variations in deposition introduced by the human element of a hand-stamp deposition process. In the case of the aforementioned and previously discussed devices, this leads to a source of high entropy, and a uniqueness of each individual device. Thus, it is expected that the use of two inks fabricated and deposited in this fashion will continue to yield unique sources of high entropy. It is also expected that, via the filtration of received optical information, each ink's emission pattern may be discerned and be considered its own individual entropy source. However, unlike comparisons between different devices, which suggest a strong sense of uniqueness that may be distilled into unique, easy to characterise identifiers; for H-PUFs it is expected that the two sources of entropy for a single device will exhibit some form of correlation — resulting in identifiers which are not necessarily completely unique. Such behaviour may stem from a variety of factors, including but not limited to, the closeness of the emission peaks (with a small amount of overlap for the full-width of the spectrum profile below the half-maximum point, as shown in Fig. 5.2); the overlap of CIS650's absorption profile and CIS530's emission profile, resulting in some "cross-talk" between the two entropy sources, as emis-

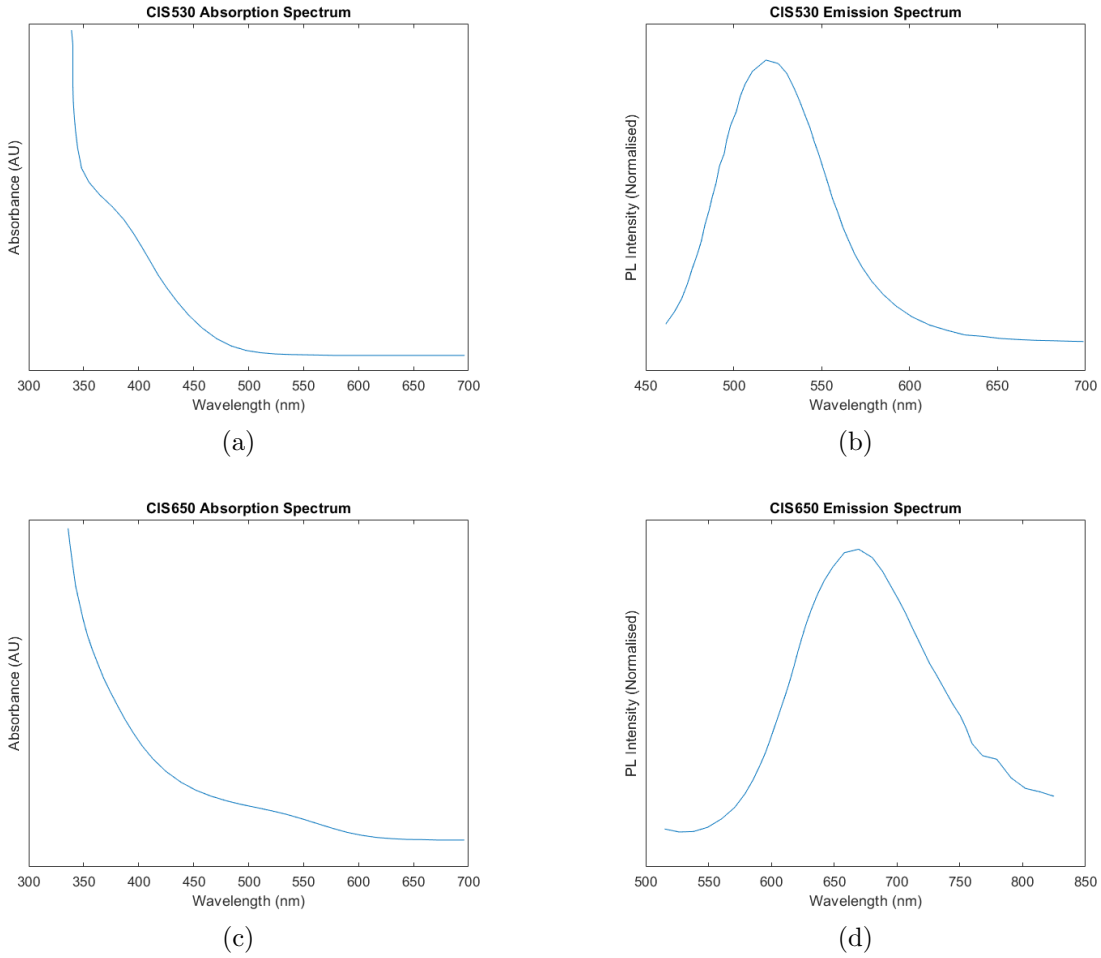

Figure 5.2: Figure showing expected emission spectra for absorbance (left) and emission (right) for CIS530 (top) and CIS650 (bottom) colloidal quantum dots used in this work. FWHM for the peak in the emission spectrum of CIS530 (CIS650) is 69 nm (100) nm. This graph was reproduced using data obtained from NNCrystal US Corporation[1].

sion of CIS530 may be absorbed by nearby CIS650 QDs, boosting intensity of the CIS650 emission where CIS530's emission is high. A simple schematic demonstrating this idea is given in figure 5.3 Additionally, and linked to the previous factors, the spatial configurations of the dot clusters themselves may contribute to any correlations; with CIS650 dots potentially being situated at the same x, y-spatial coordinates as CIS530 dots, separated only by their increased height on the surface of the device.

5.2.1 Correlation as a resource

As such, it is expected that, for a given device, the two entropy sources will exhibit a lower normalised Hamming distance between their respective generated fingerprints. At first glance, this may seem an unwanted phenomenon for a PUF: the two CRPs

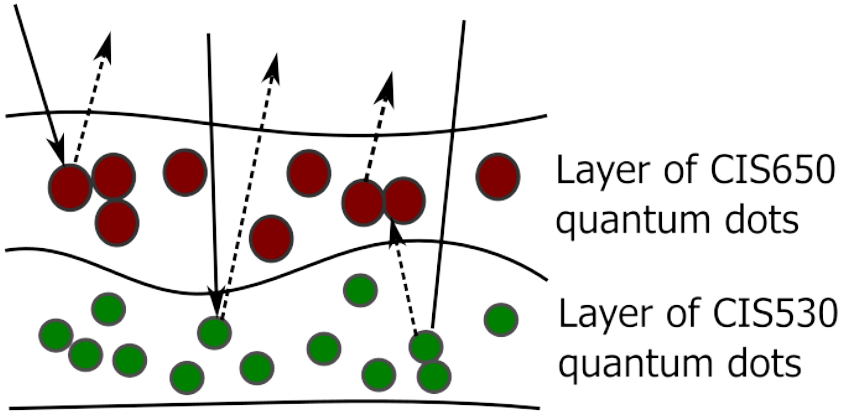

Figure 5.3: Schematic showing potential cross-talk between quantum dots in HPUFs. Incident light is represented by solid lines, whilst emission from quantum dots is represented by dashed lines.

(in reality, a single challenge, with two distinguishable responses) do not exhibit the uniqueness expected for obtaining two responses from a given device, in the manner usually expected during strong PUF analysis. However, this correlation may be used either as an additional security measure, or, potentially, as a sub-tool in a wider use-case dependent application. Here, propositions for the former is given.

5.2.1.1 Additional security via extended entropy

A simple way to make use of the second layer of quantum dots is to use it as an additional parameter for checking the authenticity of a device, for use cases such as authentification of a product. In the case of a single-dot emission device, verification is performed by ensuring that the fingerprint obtained by a user of a given device is close enough to the registered device, potentially by directly computing the Hamming distance between the user generated fingerprint x' and the registered fingerprint x; or by indirectly assessing this distance with the use of fuzzy extractors (allowing the output to be hashed for additional security in circumstances where this may be required). Given a H-PUF device instead, additional checks may be performed, making fabrication of an inauthentic device harder for counterfeiters. Here, two methods are proposed, the first involving only the two distinguishable fingerprints themselves, and the second involving the correlation also. An alternative to yield additional security is to instead combine the outputs of each response, into a single key of greater length, as done by [95]. Further, the presence of two keys allows the use of H-PUFs as reconfigurable PUFs, similar to the work by [46]. Suppose only one key, say, generated from red emission from a H-PUF was used at a time. If a communication between an honest party, Bob, in possession of a H-PUF,

and Alice, an honest verifier, were to be intercepted (and Alice and Bob detected the eavesdropping), the device need not be rendered obsolete: instead, Alice and Bob could switch to using the key obtained from green emission for future communications. By investigating the properties and uniqueness of H-PUFs with some $n > 2$ number of fine-tuned peak emission wavelengths, these could potentially be used to scale up the number of potential reconfigurations (and thus, overall challenge-response space) for a single-key PUF.

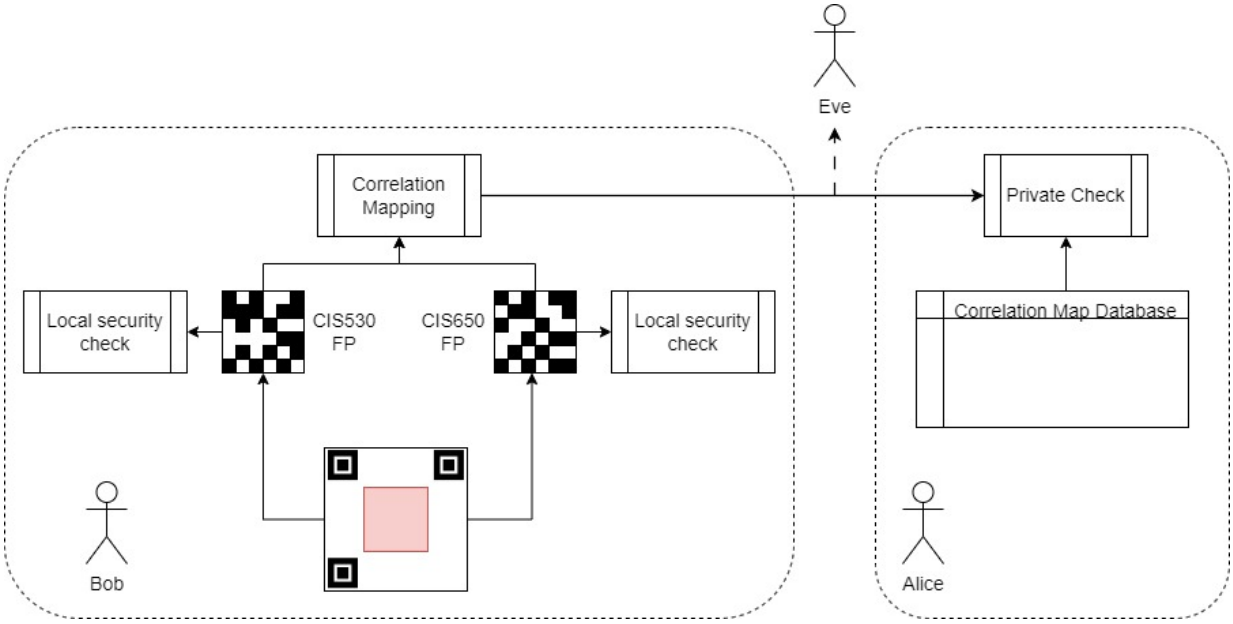

Figure 5.4: Schematic showcasing potential usecase for HPUF devices with independent correlation maps. Bob presents his PUF for local validation: two fingerprints are obtained and checked for authenticity based on closeness to a pair of registered fingerprints. If, at a later date, Alice requires knowledge of Bob's authorisation, she may request a correlation map. As long as her database is securely kept private, even if an eavesdropper, Eve, intercepts the correlation map, no other information allows her to tie it to Bob.

Additionally, correlation may be used as a method of obscuring the keys generated by each layer. In a communication protocol in which a PUF output is sent over a (necessarily) encrypted channel, correlation may be used as protection against an eavesdropper. Suppose Alice is the head of security for a high-clearance research facility, who must be able to remotely verify who is present at the facility at any given time. Additionally, due to the nature of the work and fear for the researchers' safety, their identities must not be revealed during any communications between the facility's local security system and herself. She may issue a researcher, Bob, a H-PUF as a key for the building. Upon presenting the H-PUF, the local system may interrogate it and record its responses in both channels. Upon verification that the output belongs to a trusted party, Bob is granted access to the facility.

Should Alice wish to check who is present at the facility, instead of requesting Bob's personal output keys (or, without loss of generality, hashes of the keys), which are intrinsically linked to his device, Alice may instead request a *correlation map* of the two keys. As the correlation stems from unique placement and clustering of quantum dots of various sizes, it is expected that each user's correlation map will also be unique. By checking any inbound correlation maps against her own personal, local database of employees, Alice can identify that Bob is present in the facility. If a nefarious party, Eve, were to intercept the communication between the local facility system and Alice, she would obtain no information that allows her to link the map to Bob, even if she were to have previously obtained his private keys.

Thus, the notion of one-to-many PUF is established to have at least two potential use cases. Firstly, the ability to formulate a reconfigurable PUF in the weak PUF case, using quantum dot-based inks. Secondly, it may provide an additional layer of security to use as part of the implementation of PUF-based authentication schemes, allowing for further obfuscation of protocol details, even under eavesdropping attacks.

5.3 Shroud, Full PUF Analysis

We examine the full area of the H-PUFs in an optical laboratory to ensure fine control of ambient environment, whilst allowing for easy filtration of light using camera equipment roughly analogous to the smartphone case. As outlined in Chapter 3, two effective bandpass filters are used, comprised of a combination of longpass and shortpass optical filters to capture emission chiefly from each quantum dot PUF ink individually. Example captures are shown in figure 5.5.

5.3.1 Fingerprinting Outputs

Before investigating the level of correlation between the two different responses garnered from a single device, metrics on the fingerprints for each response are presented (independent of fingerprints from a device's second response), to ensure uniqueness between devices, as well as more general requirements for a PUF. Outputs of both fingerprinting algorithms are presented in figure 5.6. Further evidence of independence of output responses is given by the reported DOF for each algorithm, shown in figure 5.8. Trends generally follow that which were established in chapter 4; with better performance for lower radii; and AHB out-performing R-MLBP.

To assess the uniqueness of devices individually, intra-and inter-Hamming distance distributions are calculated on a per-response basis: i.e., the uniqueness of a device's

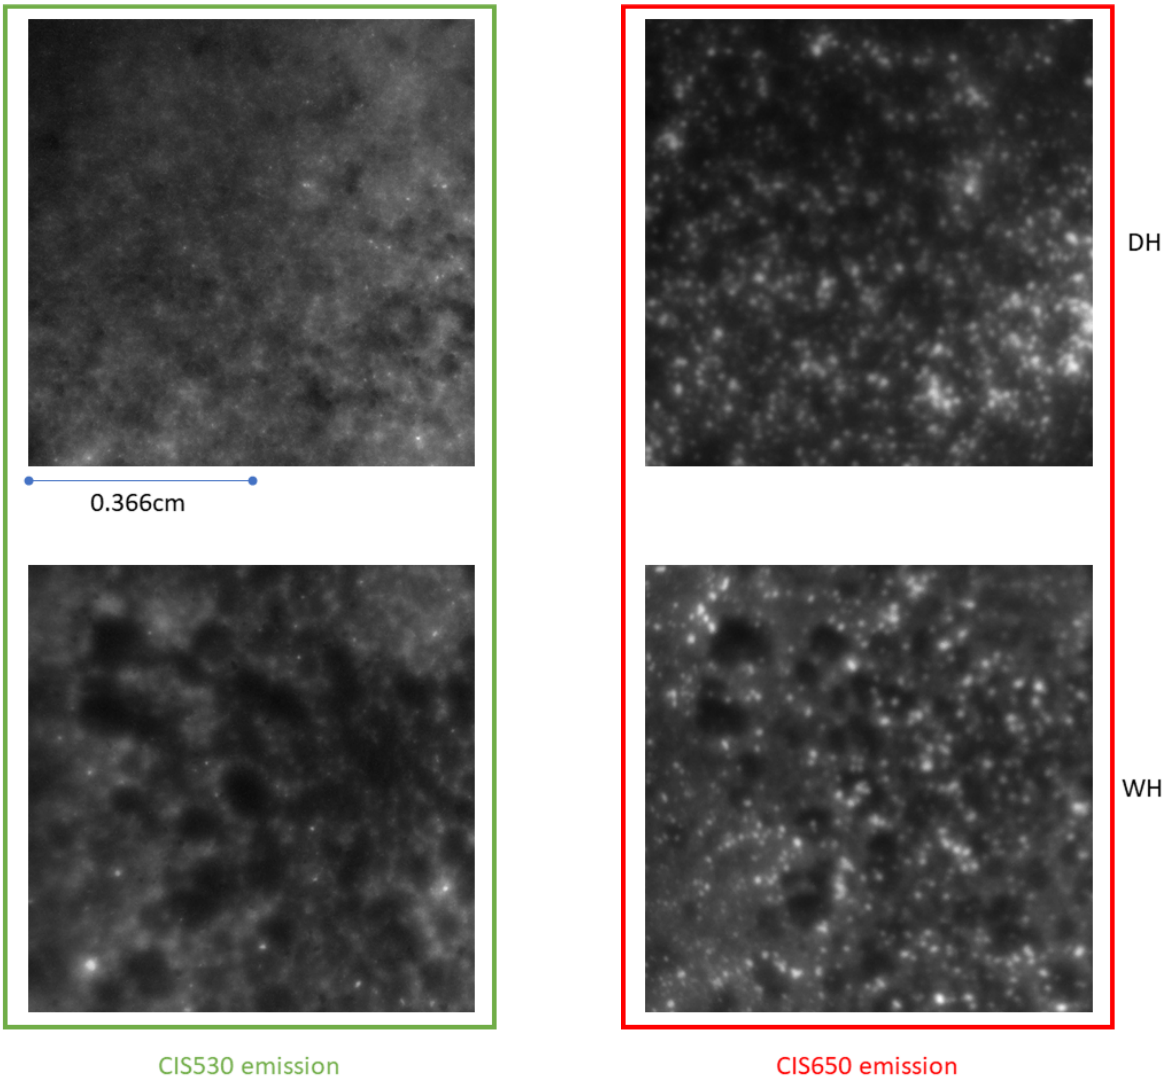

(a) (b)

Figure 5.5: Filtered images of both DH-PUF (top) and WH-PUF (bottom) devices. Images in panel (a) were acquired with a combination of 500 nm longpass and 550 nm shortpass optical filters, in order to isolate emission from CIS530 dots. Images in panel (b) were acquired with a combination of 550 nm longpass and 650 nm shortpass filters in order to isolate emission from CIS650 dots. Captures of individual responses from a single PUF appear to share some level of correlation, which can be seen in images of the WH-PUF, which share common areas with little to no emission.

CIS530 fingerprint is determined only via calculations against CIS530 fingerprints from other devices; and the reliability of the fingerprinting process will be calculated only for the set of images taken of isolated CIS530 emission. Figure 5.7 shows an example of a range of distributions for a given device. Inter-Hamming distribution means centred around 0.5 show that each device has a unique fingerprint per isolated emission output,

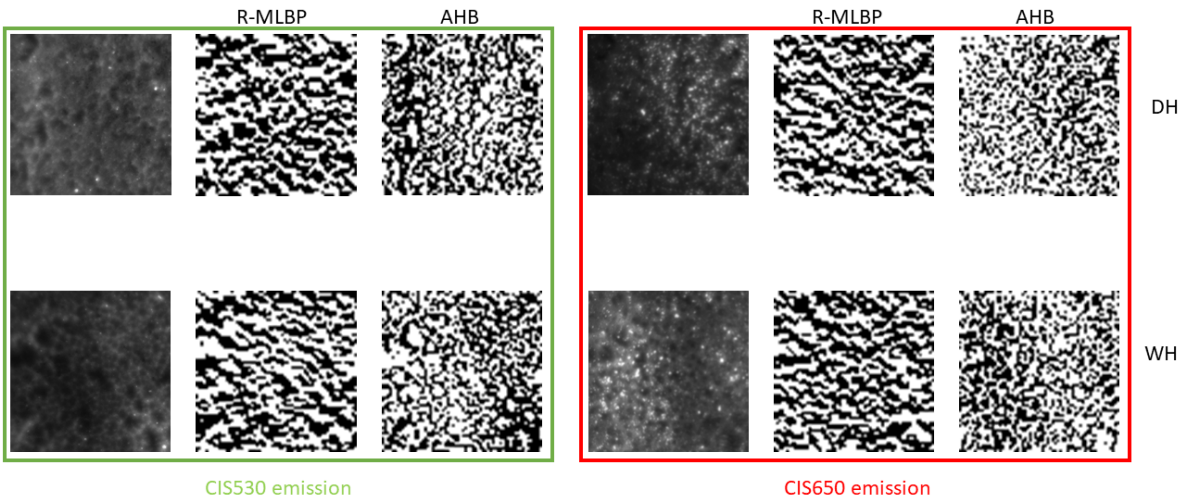

Figure 5.6: Examples of captured isolated emission from HPUF devices (isolated CIS530 on the left, with CIS650 on the right), along with their corresponding fingerprints under both R-MLBP and AHB. The top row shows examples for a DH-PUF device, whilst the bottom shows a WH-PUF device. Each greyscale PUF image constitutes a $\sim 0.732 \times 0.732$ cm^2 region.

whilst both R-MLBP and AHB are shown to be capable of reliably extracting the same fingerprint from a response by the low Intra-Hamming results.

Figure 5.8 provides an overview of how bias and DOF both vary over radius for HPUFs. Accumulatively across the entire population of fingerprints from both sets of responses from each device, bias scores for R-MLBP remain suitably close to the ideal score of 0.5, only deteriorating towards the highest choices of radii. AHB, on the other hand, once again produces worse bias figures for low choices of radius, with performance improving as radius increases. This further suggests a lack of suitability towards quantum dot OPUFs made in this manor. In the case of DOF, both algorithms perform satisfactorily — performance is similar to that seen in the case of single dot emission devices analysed in chapter 4. For R-MLBP, DOF remains above the threshold for low choices of radius, and drops to roughly the threshold for higher choices; whilst AHB tends to remain higher overall.

5.3.2 Uniqueness and Correlation

The calculated inter-Hamming distance distributions once again suggest that the uniqueness of each device is capturable via the use of both fingerprinting methods, with all devices producing fingerprints that are seemingly random when compared with each other. Paired with the low intra-Hamming distances calculated for both cases of optical filtration, there is good reason to believe that, via filtration, each individual dots' emission is

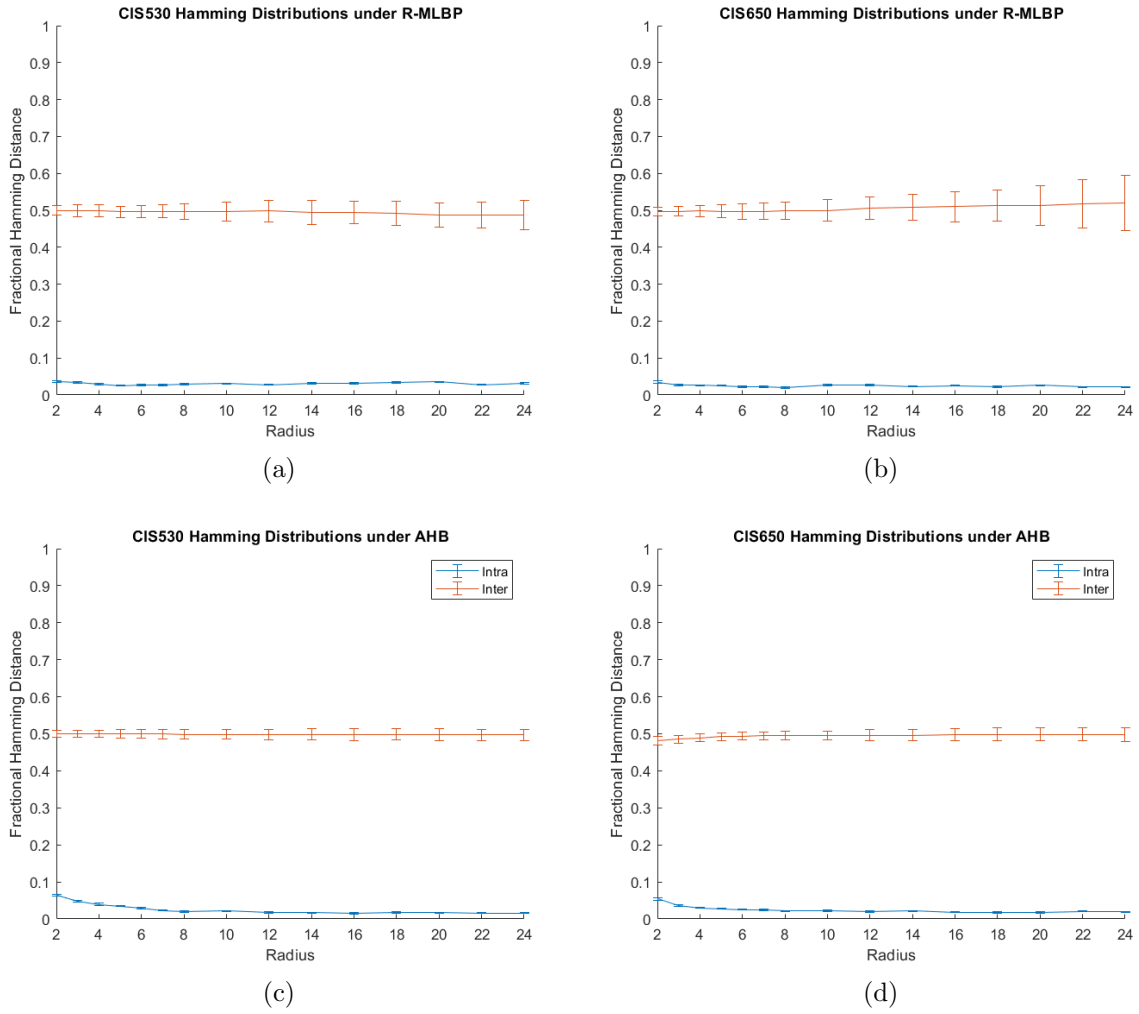

Figure 5.7: Figures showcasing intra- and inter- distributions over varying choice of radius for fingeprints generated from images of isolated emissions for a given device. Intra-distributions are based on fingerprints generated from images of a single isolated emission for a single given tag. Inter-distributons are based on fingerprints generated from images of all devices, for a single isolated emission. Distributions shown are for (a) Isolated emission from CIS530 fingerprinted under LBP; (b) Isolated emission from CIS650 fingerprinted under LBP; (c) Isolated emission from CIS530 fingerprinted under AHB; (d) Isolated emission from CIS650 fingerprinted under AHB.

reliably and uniquely capturable using the techniques discussed.

Of more interest is the Hamming distances calculated for comparisons of photos of the same device under different filtration methods. At all choices of algorithmic radius, this Hamming distance distribution (henceforth referred to as the filtered Hammings) tends to fall near the calculated inter-Hammings, often with only small separation. An example of intra-Hamming, inter-Hamming, and filtered Hamming distance for CIS650 samples is shown in figure 5.9. For both algorithms, the average of the absolute distance

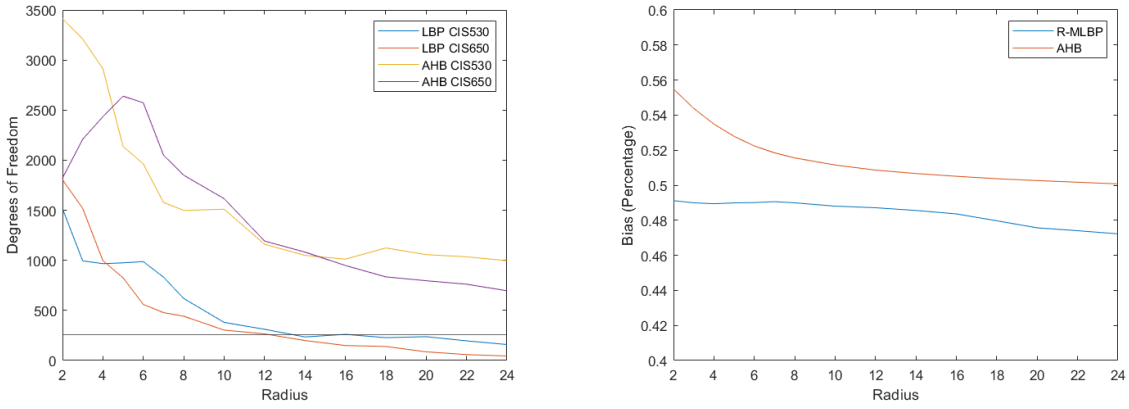

(a) (b)

Figure 5.8: Variation of (a) DOF and (b) bias for a given HPUF device. For DOF, calculations are separated for each isolated emission, and a solid line is drawn at 256 bits, marking the minimum sought-after DOF result. For bias, calculations are made using fingerprints generated for both isolated emissions under a given fingerprinting algorithm.

between filtered Hammings and inter-Hammings across all devices are 0.08 (0.09) for R-MLBP (AHB). Due to the tight standard deviations produced for both distributions, they are separable despite the relative closeness of means. However, the distributions generally lack consistency, suggesting an inconsistency to control correlation across the devices. This may limit their use in scenarios where a certain degree of correlation must necessarily exist for use, although it may be useful for protocols that wish to exploit the unpredictability of correlation levels.

One general trend noticeable in the results presented is that, as the choice of algorithmic radius, r, increases, the ability to separate the fingerprints produced by the individual dots' emission deviates further from 0.5. Noting that a normalised Hamming distance of 1 signals perfect anti-correlation, a Hamming distance $x > 0.5$ can be considered to contain the same amount of information about the fingerprint its been compared to as a Hamming distance of $y = 1 - x$. As seen in Fig. 5.5, there is a clear overlap in areas of bright emission between the two dots. Therefore, it is hypothesized that, at low choices of r, the finer, microscale differences (which may not be immediately clear by eye) in random distribution of the respective quantum dots is, to a limited degree, captured by the fingerprinting processes, whereas, as r increases, the common, macroscale features are fingerprinted, resulting in an increasing chance of information leakage (and, in a scenario where fingerprints are correlated as opposed to anti-correlated, a direct collision).

From figure 5.5, it is noticeable that a great deal of the emission pattern produced by the CIS530 QDs is present in the filtered image of the CIS650 QDs. It is hypothesised that this is the cause for the correlation that does exist, although it is not understood

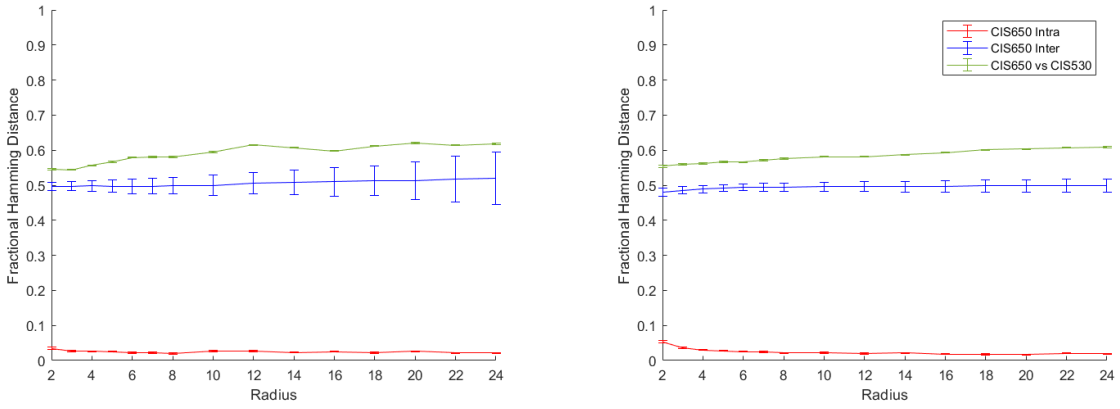

(a) (b)

Figure 5.9: Examples of the varying statistics for Hamming distributions over choices of radius for (a) R-MLBP and (b) AHB. Here, intra-(red) and inter-(blue) distributions refer to the standard intra- and inter-distributions for isolated CIS650 emission for a given tag. In green is the average calculated Hamming distance between fingerprints of the same device, under different emission filtrations (i.e., fingerprints of isolated CIS650 emission compared to fingerprints of isolated CIS530 emission.)

why this pattern is present in both images. Potentially this could be the result of the fabrication process, with ink of the second layer of dots forming around any ridges or troughs created by the uneven distribution of pressure from the handstamp process. Figure 5.11 shows an example of correlated emission patterns; in panel (a) it can easily be observed that there are areas dominated by large green (red) clusters due to the CIS530 (CIS650) dots. However, at a smaller scale, relatively minuscule clusters also form (as can be seen in the microscope images shown later in fig. 5.10), which will contribute to the optical information processed via a camera's sensor, creating correlations due to the aforementioned spatial configurations, as well as potential cross-talk between the dot types. In order to investigate this further, devices were imaged under a microscope, with the behaviour of dot clusters compared to phenomena seen in the full-device captures analysed here.

5.4 Optical Microscope Analysis

Images of H-PUFs taken under a microscope are presented, for the investigation of the root cause of shared emission patterns in full shroud imaging. the case of each image area, an optical microscope is used as described in chapter 3. Under 10 X magnification, full colour RGB images are taken at high resolution with 16-bit depth, showcasing individual clustering of quantum dots as shown in figure 5.10.

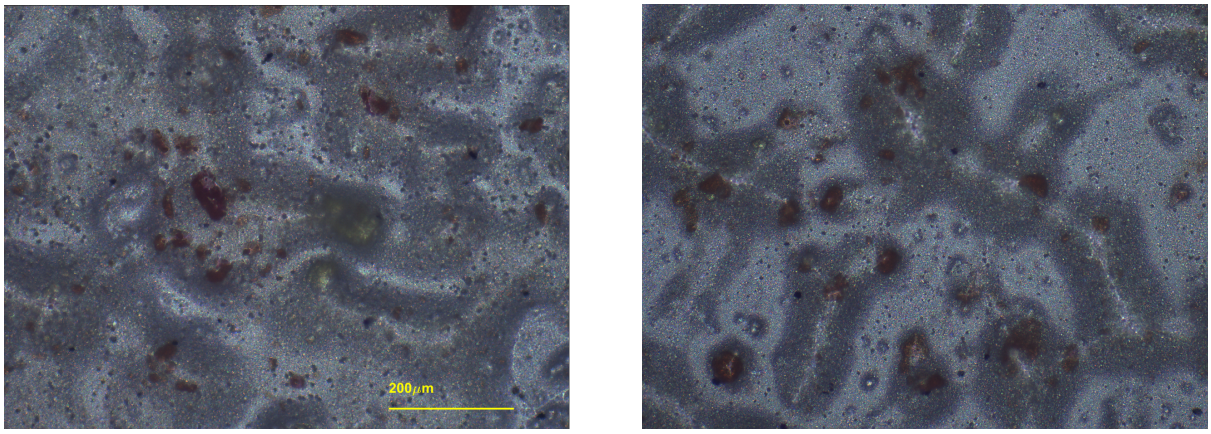

Figure 5.10: Left: Optical microscope image of a DHPUF device, with 10x magnification. Right: Optical microscope image of a WHPUF device, with 10x magnification. In both images, the tendency for CIS650 quantum dots to form larger, solid clusters can be seen.

Examining the images in Figure 5.10, it is apparent that the clustering of quantum dots within lacquer is majorly dependent on particle size, with the CIS650 quantum dots tending to conglomerate in larger, distinct clusters in comparison to the CIS530 dots, which appear to spread more diffusely through the lacquer. It is conjectured that this variation in clustering behaviour is the root of the different types of optical emission pattern present in full-shroud images: CIS650 dots tend to emit at (comparably) high levels of intensity in small discrete areas (referred to as a *constellation-like* pattern, whilst CIS530 dots tend to emit at lower levels of intensity, spread more diffusely across the device surface (referred to as a *nebulous* pattern). Returning to the images taken uner the optical microscope, it is observed that there is a lack of consistency in the deposition of lacquer across the device surface: away from clear clusters of quantum dots, the surface exhibits a mix of darker-grey areas, and light/white areas. It is conjectured that these darker grey areas, which quantum dot clusters tend to be surrounded by, are areas in which the deposition of lacquer is thicker than surrounding areas, resulting in less reflected light reaching the camera sensor. Such variations are expected to some degree, due to the uncontrollable element of applied pressure during manufacturing via the handstamping method, resulting in an uneven distribution of lacquer deposition.

Assuming such conjecture, a possible explanation for the spatial correlation of CIS530 and CIS650 dots on the devices is considered. Here, it is proposed that the drying of lacquer in varying thicknesses is the leading cause for correlation, with cross-talk between quantum dots an additional factor. Consider the different areas of full-shroud images highlighted in figure 5.11. Areas such as (a), where the CIS530 emission pattern is clearly visible in the CIS650 isolated images, with strong constellation-like clustering above, are believed to be areas in which both layers of ink have dried with thicker areas of lacquer in

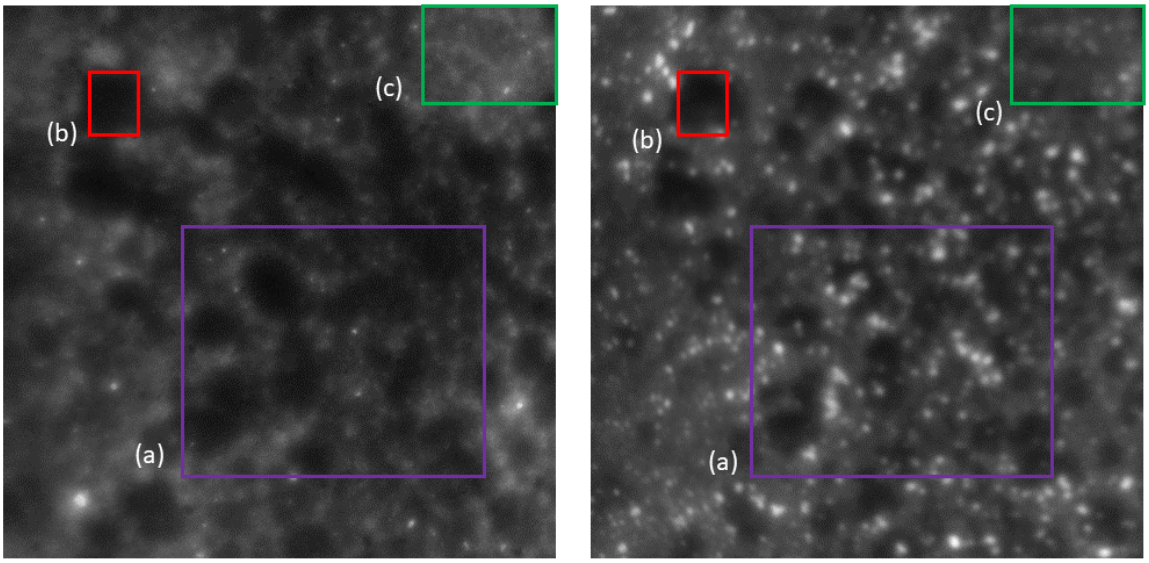

Figure 5.11: Image of a WHPUF device captured with no ambient light, within a shroud, with short-pass and long-pass optical filters fixed to the camera, allowing for isolation of emission from individual dot types. Left: Capture of isolated emission from CIS530 quantum dots. Right: Capture of isolated emission from CIS650 quantum dots. Highlighted areas: (a) Purple (bottom-most); (b) Red (top left-most); (c) Green (top right-most). Each greyscale PUF image constitutes a $\sim 0.732 \times 0.732$ cm^2 region.

the same location, and a strong clustering of CIS650 dots. Meanwhile, areas such as (b), in which little emission is present in both filtered images, are suggested to stem from a thin deposition of lacquer from both inks. Areas such as (c), where the CIS530 emission pattern is clearly visible in the CIS650 isolated images, but there is no overlying strong constellation-like pattern, are posited to be areas in which either: (i) thick areas of lacquer have formed from deposition of CIS530, whilst the deposition of CIS650 only contributed a thin layer of lacquer; or, (ii), thick areas of lacquer have formed from deposition of both CIS530 and CIS650 inks, yet CIS650 has not formed a high concentration of clusters in this area. For the case where sections of a device are characterised by (c), we believe that the presence of the emission pattern may be boosted by cross-talk between the two types of dots: areas with a strong emission pattern from CIS530, but only a thin, diffuse layer of CIS650 appear similar to the isolated CIS530 image due to the emission from CIS530 further exciting the thin layer of CIS650 and increasing the emission intensity from these dots.

5.5 Concluding Remarks

The HPUF tags presented in this chapter are confirmed to be strong sources of static entropy of optical information centred around both 530 nm and 650 nm wavelength emis-

sions, for which outputs of one device are highly independent from outputs of another device, similarly to the single-wavelength emission devices examined in the earlier chapter 4. Further, information obtained from the two distinct entropy sources is shown to be intrinsically linked, with Hamming distances between fingerprints of different isolated emission wavelengths of the same tag shown to typically fall slightly above or below 0.5, which would signal comparative uniqueness. This shared information takes the form of an increased correlation when the associated Hamming distance is under 0.5; and anti-correlation when the associated Hamming distance is above 0.5. Areas of correlation and anti-correlation within fingerprints of each independently isolated emission are believed to be highly related to the irregularities in lacquer deposition present in the manufacturing method. Proposals are put forward for how such devices, with varying levels of correlation, may be exploited for use in communication protocols, viable with current technology.

Chapter 6

Mixed State Compilation

The task of learning an unknown quantum state is a fundamental primitive in quantum computing, and one that can be approached in a variety of ways. One method, which has been studied extensively, is *quantum tomography* (a review of tomography techniques may be found here [25]) — the act of learning the matrix elements of an unknown state, ρ, directly by making a series of (known) measurements on multiple samples of the (unknown) state. By combining the results of such measurements and making use of Born's rule, one ends up with a full, classical description of the state. It has been conjectured [117] that using only $\Omega(\text{rank}(\rho)\text{d}/\epsilon^2)$ measurements is sufficient to gain a description of the density matrix of ρ, up to additive error ϵ in the trace distance. Obtaining a full, classical description of a state by such means naturally requires the number of measurements to increase exponentially as the number of qubits grows. As a result, for large systems quantum state tomography falls victim to *the curse of dimensionality*, a phrase coined by Bellman [11], encapsulating the phenomenon of sparse, high-dimensional data leading to the intractability of learning problems. Further research has led to the development of "specialised" tomography-based tasks, such as matrix product state tomography [24, 119] and neural network tomography [107]; where, given a state with certain known properties, the required number of state samples is reduced. However, for unknown systems whose properties are not known (or, for some reason, can not easily be found out), a general system for learning a state of arbitrary size remains evasive. As opposed to seeking a full description of ρ's density matrix, we can instead seek to learn some operational property of ρ in a way that is sufficient to characterise the state, either fully, or for specific properties required for a task at hand. This avenue was first explored by Aaronson in [2], leading to the task of *shadow tomography*. Shadow tomography exploits the linearity of many interesting quantum properties' function on the density matrix of ρ. It has been found that learning the output of such functions (and hence, a description of that property of a state) can require significantly fewer copies of the state,

avoiding the curse of dimensionality (or, at the very least, putting it on hold). However, whilst traditional shadow tomography reduced the number of samples of ρ required for learning, it employed quantum circuits with vastly great depth, and required working quantum memory, making it impractical on today's NISQ hardware. Algorithms reducing this quantum burden have since been proposed, often by introducing techniques from (classical) machine learning, such as semidefinite programming.

An alternative approach to learning an unknown state is the task of *quantum state compilation*, in which one learns the quantum circuitry required to prepare a state. For a known state, compilation can prove useful in learning a hardware-specific, efficient circuit with which to prepare the state on quantum hardware for further processing[1]. Compilation may also be used a means for compression of a quantum state, in which one wishes to find an approximation of a full (potentially unknown) state that can be stored with fewer resources (qubits) than required for a full description of the state. In the case of prioritising a select number of eigenvalues of the state for its representation, compression via compilation naturally extends to the task of *principal component analysis* (PCA). Applications of compilation may also extend to work on PUFs whose challenges and responses take the form of quantum states. Compilation may help in the preparation of challenge states, by providing more hardware efficient circuits for the task; as well as directly learning the output of a quantum PUF.

The rest of this chapter is organised as follows. First, a variation quantum algorithm for mixed state compilation is defined, with a particular emphasis on the operational formulation of the ansätze used in evaluating the cost function, the Hilbert-Schmidt distance. Then, it is briefly shown that the chosen cost function's gradient is analytically computable. The rest of the chapter deals with alternative formulations of the cost function, constructing local alternatives to the Hilbert-Schmidt distance in order to increase trainability of the algorithm. Each formulation is analysed for its suitability as a cost function in different use cases.

6.1 Quantum Mixed State Compiling

In [39], a variational quantum algorithm for state compilation, dubbed the *Quantum Mixed State Compiling Algorithm* (QMSC) was introduced. Whilst the focus of this chapter is principally on analysis of the cost function (and variants of it) used in QMSC, we will first define the algorithm, in order to highlight specific use cases of it which will be important in the later analysis.

[1] Or, in the case that such a circuit is already known, compilation may prove useful in learning a more efficient, lower-depth circuit.

6.1.1 QMSC Algorithm

The QMSC algorithm takes as its input a (typically unknown) target state ρ, along with a desired approximation rank, R, and outputs a trial state, σ, with rank R. If the desired approximation rank is chosen such that $R \geq \text{rank}(\rho)$, a successful implementation will result in an output $\sigma \approx \rho$. If instead, one seeks to learn a lower rank approximation of ρ, (i.e., $R < \text{rank}(\rho)$, a successful implementation will result in an output $\sigma \approx \rho^*$, where ρ^* is the solution to the quantum low-rank approximation problem (QLRAP), given in [40]. The degree to which the output σ is "close" to ρ is given in terms of the cost function chosen for QMSC, the *Hilbert-Schmidt Distance* between ρ and σ.

Definition 6.1. The Hilbert-Schmidt (HS) distance is a distance measure between two operators in Hilbert space,

$$D_{\text{HS}}(\rho, \sigma) = \|\rho - \sigma\|_2^2 \equiv \text{Tr}\left[\rho^2\right] + \text{Tr}\left[\sigma^2\right] - 2\text{Tr}\left[\rho\sigma\right], \tag{6.2}$$

where $\|\cdot\|_2^2$ denotes the *Schatten* 2*-norm*, as introduced in chapter 2.

An analysis of the suitability of the HS distance as a cost function is provided in Sec. 6.2.

For the preparation of our trial state, we deploy two different ansätze, for reasons which will briefly be touched upon later (for a fuller description of the importance of the ansatz choice, refer to [39]). The ansätze are defined as follows:

Definition 6.3. For the input of the trial state, σ, the *Convex Combination of Pure States* (CCPS) ansatz takes the form

$$\sigma_{\text{CCPS}}(\vec{\theta}, \vec{\phi}, R) := \sum_{i=0}^{R-1} p_{\vec{\phi}}(i) U_{\vec{\theta}} \left|i\right\rangle \left\langle i\right| U_{\vec{\theta}}^{\dagger}, \tag{6.4}$$

where $\{|i\rangle\}_{i=0}^{R-1}$ denotes a subset of the computational basis on n qubits, $\vec{\theta}, \vec{\phi}$ are vectorised parameters, $U_{\vec{\theta}}$ is a quantum circuit parametrised by $\vec{\theta}$, and $p_{\vec{\phi}}$ is a probability distribution parametrised by $\vec{\phi}$.

Definition 6.5. For the input of σ, the *State Purification* (SP) ansatz takes the form

$$\sigma_{\text{SP}}(\vec{\theta}, R) := \text{Tr}_{\text{A}}\left[\text{U}_{\vec{\theta}} \left(\left|0\right\rangle \left\langle 0\right|\right)^{\otimes(\text{n}+\text{n}_{\text{A}})} \text{U}_{\vec{\theta}}^{\dagger}\right], \tag{6.6}$$

where $U_{\vec{\theta}}$ is a parametrised quantum circuit acting on the n system qubits, and an additional n_A ancilla qubits, where n_A is determined by the choice of R.

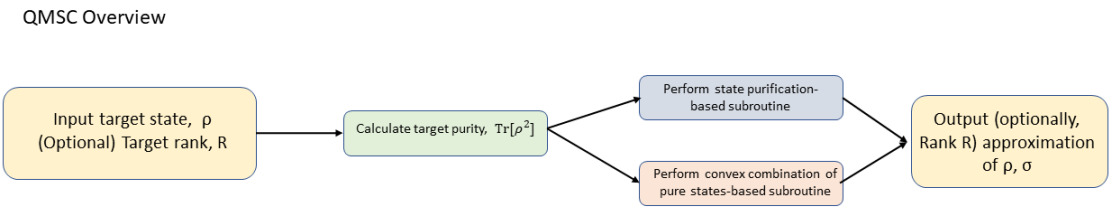

Figure 6.1: Simple flow diagram depicting most general overview of QMSC algorithm.

Where context leaves no room for ambiguity, we will opt to drop the CCPS and SP subscripts.

The general form of the QMSC algorithm is represented in Fig. 6.2, and formally described as follows:

1. Input target state, ρ, along with desired output rank R for the output trial state, σ.

2. Calculate target purity term, $\mathrm{Tr}\left[\rho^2\right]$.

3. Perform optimisation subroutine as follows:

   - For CCPS ansatz:
     (a) Evaluate overlap terms via Loschmidt echo.
     (b) Compute cost function $\mathrm{Tr}\left[\rho^2\right] + \sum_{\mathrm{i}} \mathrm{p}_{\tilde{\phi}}(\mathrm{i}) - 2\sum_{\mathrm{i}} \mathrm{p}_{\tilde{\phi}}(\mathrm{i}) \left\langle \mathrm{i}\right| \mathrm{U}_{\tilde{\theta}}^{\dagger} \rho \mathrm{U}_{\tilde{\theta}} \left|\mathrm{i}\right\rangle$.
     (c) Update parameters $\vec{\theta}$ and $\vec{\phi}$ via classical optimisation, and return to step (a).
   - For SP Ansatz:
     (a) Evaluate overlap term and ansatz purity via Loschmidt echo or SWAP circuits.
     (b) Compute cost function $\mathrm{Tr}\left[\rho^2\right] + \mathrm{Tr}\left[\sigma^2\right] - 2\mathrm{Tr}\left[\rho\sigma\right]$.
     (c) Update parameter $\vec{\theta}$ via classical optimisation, and return to step (a).

4. Output final trial state σ^{OUT} of rank R.

Upon successful implementation of the algorithm, both ansätze form a trial state $\sigma \equiv \rho$ in the case of $R = \mathrm{rank}(\rho)$, and $\sigma \equiv \rho^*$ for $R < \mathrm{rank}(\rho)$. For the SP ansatz, we may only learn lower-rank approximations where R is a power of two, due to the link between R and the number of ancillary qubits, n_A. We note that this does not limit the ability to learn full-rank approximations of a target state, as choosing $2^{n_A} > \mathrm{rank}(\rho)$ will yield a purification of suitable rank upon discarding the ancillary system. This is

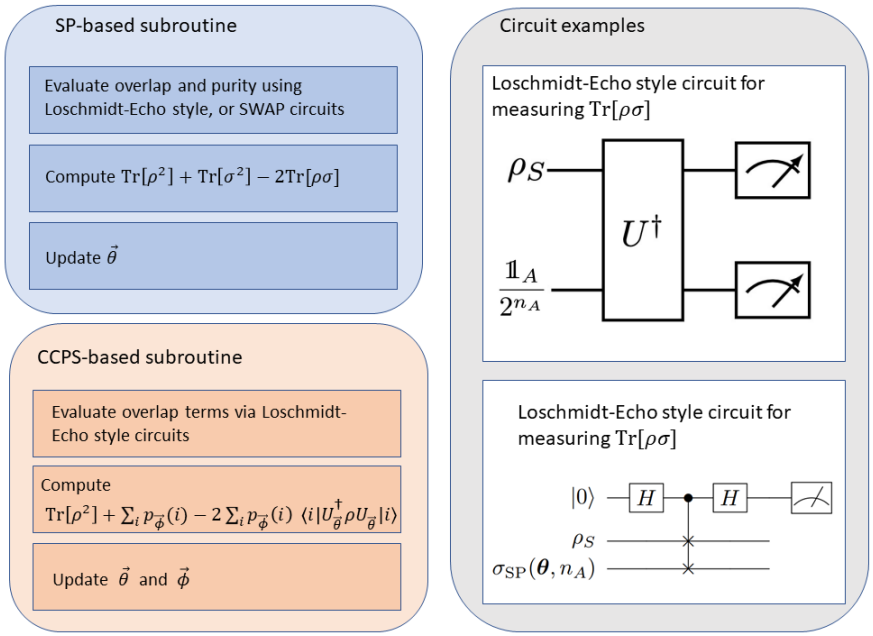

Figure 6.2: Oveview of both the state-purification (left-upper), and convex combination of pure states-based (left-lower) subroutines and examples of circuitry that may be used to evaluate terms (right).

in contrast to the CCPS ansatz, for which there is no ancillary system and instead R can be chosen freely. In the case of the CCPS ansatz however, one must perform an additional optimisation step, to account for the trainable terms contained in the classical description of the probability vector, $p_{\vec{\phi}}$. We note further that, in practicality, $p_{\vec{\phi}}$ may only be stored classically for either small-scale problems, or for choices of small R: for large-scale problems and/or large R, the number of terms scales exponentially, necessitating sampling and adding a further layer of complexity to the process. For a full description of algorithmic implementation and circuits used, as well as both simulated and experimental results of performance, refer to [39].

6.1.2 Gradient Analysis

In order to successfully run the proposed VQA for mixed state compilation using gradient-based optimisers, we must be able to compute an analytic expression for our cost function's gradient, such that it can be calculated at each necessary step. Here, we provide a brief calculation of the analytical gradient for the Hilbert-Schmidt distance.

6.1.2.1 Global Cost

Firstly examining the state purification ansatz, we recall the appropriate Hilbert-Schmidt distance cost,

$$C_{\mathrm{SP}(\vec{\theta},n_A)} = \mathrm{Tr}\left[\rho^2\right] + \mathrm{Tr}\left[\sigma(\tilde{\theta},\mathrm{n_A})^2\right] - 2\mathrm{Tr}\left[\rho\sigma(\tilde{\theta},\mathrm{n_A})\right], \tag{6.7}$$

where our trial state, $\sigma(\vec{\theta},n_A)$ is typically expressed via its Stinespring representation, i.e.,

$$\sigma(\vec{\theta},n_A) := \mathrm{Tr_A}\left[\mathrm{U}_{\tilde{\theta}}\left(|0\rangle\langle 0|\right)^{\otimes(\mathrm{n}+\mathrm{n_A})}\mathrm{U}_{\tilde{\theta}}^{\dagger}\right]. \tag{6.8}$$

Computing the gradient w.r.t. $\vec{\theta}$, we see that

$$\begin{aligned}\nabla C_{\mathrm{SP}(\vec{\theta},n_A)} &= \nabla\mathrm{Tr}[\rho^2] + \nabla\mathrm{Tr}\left[\sigma(\tilde{\theta},\mathrm{n_A})^2\right] - 2\nabla\mathrm{Tr}\left[\rho\sigma(\tilde{\theta},\mathrm{n_A})\right] \\ &= \nabla\mathrm{Tr}\left[\sigma(\tilde{\theta},\mathrm{n_A})^2\right] - 2\nabla\mathrm{Tr}\left[\rho\sigma(\tilde{\theta},\mathrm{n_A})\right]\end{aligned} \tag{6.9}$$

Due to the lack of dependence upon $\vec{\theta}$ in the purity term for ρ. The overlap term obeys the standard Pauli parameter shift rule[2], in which we implement a shift on the circuitry of $\theta \mapsto \theta \pm \frac{\pi}{2}$. The gradient of the purity term for our trial state $\sigma(\vec{\theta},n_A)$ is less trivial, owing to the squaring of our parameter-dependent term. However, exploiting the fact that, for some function $f(x)$, we have that

$$\frac{\partial}{\partial\theta}\mathrm{Tr}\left[\mathrm{f(A}(\theta))\right] = \mathrm{Tr}\left[\mathrm{f'(A}(\theta))\frac{\partial\mathrm{A}}{\partial\theta}\right]. \tag{6.10}$$

We can rewrite (individual terms of) this term as

$$\frac{\partial\mathrm{Tr}\left[\sigma(\tilde{\theta},\mathrm{n_A})^2\right]}{\partial\theta_k} = 2\mathrm{Tr}\left[\sigma\left(\tilde{\theta},\mathrm{n_A}\right)\frac{\partial\sigma(\tilde{\theta},\mathrm{n_A})}{\partial\theta_{\mathrm{k}}}\right], \tag{6.11}$$

from where we can again apply the Pauli parameter shift rule, applied such that

$$\frac{\partial\mathrm{Tr}\left[\sigma(\tilde{\theta},\mathrm{n_A})^2\right]}{\partial\theta_k} = \mathrm{Tr}\left[\sigma(\tilde{\theta},\mathrm{n_A})\sigma\left(\tilde{\theta}+\frac{\pi}{2}\cdot\tilde{\mathrm{e}}_{\mathrm{k}},\mathrm{n_A}\right)\right] - \mathrm{Tr}\left[\sigma(\tilde{\theta},\mathrm{n_A})\sigma\left(\tilde{\theta}-\frac{\pi}{2}\cdot\tilde{\mathrm{e}}_{\mathrm{k}},\mathrm{n_A}\right)\right], \tag{6.12}$$

where $\vec{e_k}$ is a standard basis vector. As such, we have all requisite information to calculate the analytical expression for the gradient of our cost function. For the local costs defined

[2]Given a variable to be computed via quantum machinery, and its gradient, if the gradient may be expressed in the original circuitry, but under a different parametrisation, we may compute it simply by shifting the parameter in our circuit.

below, the gradient analysis is entirely analogous, and is thus omitted from this thesis for brevity.

6.2 Search for a local cost function

In [39], it is shown that the above algorithm successfully compiles quantum states and their low-rank approximations for small-scale problems, with better performance shown for selected well-defined classes of states. However, due to the known existence of barren plateaus in VQA cost landscapes, it is expected that trainability will be greatly reduced as the problem size scales. As such, in a similar vein to the work shown in [69], we may wonder if a well-defined local cost may be found for the Hilbert-Schmidt distance, yielding a cost-landscape less susceptible to barren plateaus.

For a cost function to be suitable for training, we desire that it establishes the following criteria: An appropriate cost function should be:

1. Faithful — The cost function should vanish iff the compilation is exact, i.e., $\sigma = \rho$;
2. Efficiently computable on quantum hardware
3. Scaleable — As the problem size grows, complexity of compilation should grow linearly with it;
4. Operationally meaningful — non-zero values should provide some physical or operational description about the objects in question.

We note that in the case of seeking a local alternative to a faithful global cost function, faithfulness is equivalent to the requirement that $C_{\text{Local}} = 0 \iff C_{\text{Global}} = 0$.

Recall our (ansatz-blind) global cost function, the Hilbert-Schmidt distance:

$$\text{HS}(\rho, \sigma) = \text{Tr}[\rho^2] + \text{Tr}[\sigma^2] - 2\text{Tr}[\rho\sigma], \tag{6.13}$$

which comprises of the purity of each state (first two terms) and the overlap between the two (the final term). This cost is (almost trivially) faithful, and it naturally has an operational meaning based on the above, as well as the (more mathematically abstract) consideration as the Schatten 2-norm on $\rho - \sigma$. It can be calculated efficiently on quantum hardware using the either the SWAP test or a Loschmidt-Echo style circuit, as defined in chapter 2. However, due to the globality of the cost, it falls victim to the curse of barren plateaus, creating the need to establish a local-version of the cost. Two frameworks for local costs are explored in this chapter: the use of locally defined inputs to the algorithm; and the use of locally defined operators within the ansätze preparation.

6.2.1 Local States-Based Attempts

Perhaps the simplest way to formulate a local version of the Hilbert-Schmidt distance cost is to invoke the Hilbert-Schmidt distance as is, but considering marginal states on ρ and σ as opposed to the full states. For a complete description on all n qubits, this would involve making n local assessments of the Hilbert-Schmidt distance, and taking the average.

Definition 6.14. For a given target state, ρ, and a trial state, σ, we define the *Marginal* 1*-local cost* as

$$C_{\mathrm{M}}^{(1)} := \frac{1}{n}\sum_{j=1}^{n} \|\rho_j - \sigma_j\|_2^2 \,. \tag{6.15}$$

where $\rho_j = \mathrm{Tr}_{\bar{\mathrm{j}}}[\rho]$ and $\sigma_j = \mathrm{Tr}_{\bar{\mathrm{j}}}[\sigma]$, where $\bar{j}$ denotes the set of all system qubits other than the j-th qubit.

The marginal 1-local cost is naturally an operational meaningful cost — it retains the same meaning as the global cost, except it is descriptive of marginals as opposed to full states. Likewise, each term in the cost can be efficiently computed on quantum hardware in the same manner as the global cost (although the number of shots required will naturally scale up with the number of marginals being compared). However, the faithfulness of the global cost does not translate to the local case so trivially. The marginal 1-local cost is only faithful in certain cases, specifically for tensor product states, as will be shown next.

Proposition 6.16. *The marginal* 1*-local cost is faithful for tensor product states.*

Proof. For tensor product states, ρ and σ take the form $\rho = \bigotimes_{j=1}^{n} \rho_j$ and $\bigotimes_{j=1}^{n} \sigma_j$. Our cost function becomes

$$\begin{aligned} C_{\mathrm{M}}^{(1)}(\rho,\sigma) &= \frac{1}{n}\left(\|\mathrm{Tr}_{\bar{1}}[\rho] - \mathrm{Tr}_{\bar{1}}[\sigma]\|_2^2 + \|\mathrm{Tr}_{\bar{2}}[\rho] - \mathrm{Tr}_{\bar{2}}[\sigma]\|_2^2 + \cdots + \|\mathrm{Tr}_{\bar{n}}[\rho] - \mathrm{Tr}_{\bar{n}}[\sigma]\|_2^2\right) \\ &= \frac{1}{n}\left(\|\rho_1 - \sigma_1\|_2^2 + \|\rho_2 - \sigma_2\|_2^2 + \cdots + \|\rho_n - \sigma_n\|_2^2\right). \end{aligned} \tag{6.17}$$

As each term $\|\rho_i - \sigma_i\|_2^2 \geq 0$ with equality iff $\rho_i = \sigma_i$, if any individual elements of the tensor product representation of ρ differ to the equivalent element of the tensor product representation of σ, then the cost will be non-zero (and of course, the states will not be the same). Trivially, if $C_{\mathrm{M}}^{(1)}(\rho,\sigma) = 0$, then $\rho_i = \sigma_i \,\forall\, i$ and thus $\rho = \sigma$. □

However, it's easy to see that this cost will not be universally faithful, due to the fact that different entangled states may have different reduced states. A simple example to showcase this can be found using orthogonal Bell states, i.e., $\rho = |\Psi_+\rangle\langle\Psi_+|$ and

$\sigma = |\Psi_-\rangle\langle\Psi_-|$, whose reduced states are all maximally mixed. Calculating the marginal 1-local cost for these states yields

$$\begin{aligned} C_{\mathrm{M}}^{(1)}(\rho,\sigma) &= \frac{1}{2}\left(\|\mathrm{Tr}_{\bar{1}}\left[|\Psi_+\rangle\langle\Psi_+|\right] - \mathrm{Tr}_{\bar{1}}\left[|\Psi_-\rangle\langle\Psi_-|\right]\|_2^2 + \|\mathrm{Tr}_{\bar{2}}\left[|\Psi_+\rangle\langle\Psi_+|\right] - \mathrm{Tr}_{\bar{2}}\left[|\Psi_-\rangle\langle\Psi_-|\right]\|_2^2\right) \\ &= \frac{1}{2}\left(\|\mathbb{I}/2 - \mathbb{I}/2\|_2^2 + \|\mathbb{I}/2 - \mathbb{I}/2\|_2^2\right). \end{aligned} \tag{6.18}$$

Locality of a cost function need not only be defined in terms of single-qubit based operations. Instead, we can seek to explore the range of number of qubits from 1 to $k < n$ (where n is the total number of qubits in the system). By adapting $C_{\mathrm{M}}^{(1)}$ to include k-local terms, we can construct a faithful cost based on operations similar to those described above.

Definition 6.19 (Generalised marginal cost)**.** Let ρ and σ be quantum states defined on a set of n qubits, $\mathcal{Q}$. Further, group $\mathcal{Q}$ into some (potentially multiple) partitioning(s) $\mathcal{P}_k$, consisting of N_k subsets $\{\mathcal{Q}_1^{(k)}, ..., \mathcal{Q}_{N_k}^{(k)}\} \subset \mathcal{Q}$ of (at most) $k < n$ qubits such that $|\mathcal{Q}_l^{(k)}| = k$ for at least one $l \in [1, N_k]$, define the *Generalised marginal cost* as follows:

$$C_{\mathrm{M}}^{\mathrm{gen}}(\rho,\sigma) := \sum_{k=1}^{n} \alpha_k(t) C_{\mathrm{M}}^{(k,\mathcal{P}_k)}(\rho,\sigma), \tag{6.20}$$

where

$$C_{\mathrm{M}}^{(k,\mathcal{P}_k)}(\rho,\sigma) := \frac{1}{N_k}\sum_{i=1}^{N_k} \|\rho_i^{(k)} - \sigma_i^{(k)}\|_2^2, \tag{6.21}$$

with $\{\rho_i^{(k)}, \sigma_i^{(k)}\}_{i=1}^{N_k}$ forming the set of marginals on ρ and σ as determined by $\mathcal{P}_k$, and the iteration dependent weightings $\alpha_k(t) \geq 0$ satisfying $\sum_k \alpha_k(t) = 1$; and defining $C_{\mathrm{M}}^{n,\mathcal{P}_n}$ (i.e., the "n-local" term) to be our global cost.[3]

This generalised cost could, from a practical perspective, be used in an adaptive fashion over the course of training. One could start off with a weighting (described by $\alpha_k(l)$) that prioritises low-k terms, and adapt the weighting to shift towards the global cost over time (as one gets closer to the target state). As it is trivial that $\rho = \sigma \implies C_{\mathrm{M}}^{(k,\mathcal{P}_k)} = 0$ for any choice of $k, \mathcal{P}_k$, then, as long as $\alpha_k(l) \neq 0$ for $k = n$, then this cost is faithful to the global equivalent, as this global term will ensure inequality when $\rho \neq \sigma$.

6.2.2 Local Measurements on Global States

We may alternatively seek to define a local cost via *local measurements* as opposed to local states, using the light-cone argument presented in [19]. Above, a cost function

[3]Note that here we only have one choice of set $\mathcal{P}_n$, the set of all qubits.

was constructed by considering locality based on states. As such, the relevant circuitry for parametrising and computing the cost could be considered "blind" to the locality of the cost function, in the sense that the circuit could accept any state, of any ansatz type; and in this case the states in question just happened to be local (reduced) states. However, in the case of local measurements, such cost functions cannot remain ansatz-blind — that the operational formula of the cost function assessment (i.e., the circuitry) is specifically geared to the ansatz type requires local costs based on local measurements to be constructed in an ansatz-specific manner.

In order to construct such a cost, it is necessary to identify which parts of the circuit are parametrised and optimised on quantum hardware, and replace global components in these with local components. Recall the cost function as formulated for the CCPS ansatz:

$$\begin{aligned} C_{\mathrm{CCPS}}(\rho,\sigma) &:= \mathrm{Tr}[\rho^2] + \sum_{\mathrm{i}} \mathrm{p}_\phi(\mathrm{i})^2 - 2\sum_{\mathrm{i}} \mathrm{p}_\phi(\mathrm{i}) \left\langle \mathrm{i}\right| \mathrm{U}_{\tilde{\theta}}^\dagger \rho \mathrm{U}_{\tilde{\theta}} \left|\mathrm{i}\right\rangle \\ &= \mathrm{Tr}[\rho^2] + \sum_{\mathrm{i}} \mathrm{p}_\phi(\mathrm{i})^2 - 2\sum_{\mathrm{i}} \mathrm{p}_\phi(\mathrm{i}) \mathrm{Tr}_{\mathrm{S}} \left[\mathrm{U}_{\tilde{\theta}}^\dagger \rho \mathrm{U}_{\tilde{\theta}} \mathrm{H}_{\mathrm{G}}^{(\mathrm{i})}\right], \end{aligned} \tag{6.22}$$

where $H_{\mathrm{G}}^{(i)} := |i\rangle \langle i|_S$ emphasises the globality of the operations. The first term, the purity of the target state, ρ, is not optimised, so can retain its global form. The second term, the purity of the training state, σ, is optimised over. However, as outlined in [39], this optimisation process can be done purely classically, allowing us to again return the global form. This leaves the overlap term, which is optimised on quantum hardware, to be modified to include local measurements. That is, the global operator $H_{\mathrm{G}}^{(i)}$ must be replaced with some (set of) local measurement(s). There is freedom in the choosing of such measurements. One simple approach, similarly to the naïve attempt of constructing a marginal 1-local cost is to choose the average of all 1-local measurements.

Definition 6.23 (Local CCPS cost)**.** We can define a local cost function for the CCPS ansatz as

$$C_{\mathrm{L}}^{\mathrm{CCPS}} = \mathrm{Tr}\left[\rho^2\right] + \sum_{\mathrm{i}} \mathrm{p}_\phi(\mathrm{i})^2 - 2\sum_{\mathrm{i}} \mathrm{p}_\phi(\mathrm{i}) \mathrm{Tr}_{\mathrm{S}} \left[\mathrm{U}_{\tilde{\theta}}^\dagger \rho \mathrm{U}_{\tilde{\theta}} \mathrm{H}_{\mathrm{L}}^{\mathrm{CCPS(i)}}\right], \tag{6.24}$$

where we define

$$H_{\mathrm{L}}^{\mathrm{CCPS}(i)} := \frac{1}{n} \sum_{j=1}^{n} |i_j\rangle \langle i_j|_{S_j} \otimes \mathbb{I}_{S_{\bar{j}}} \tag{6.25}$$

where $|i_j\rangle$ denotes the j-th bit of the computational basis element (on n qubits) bit-string $|i\rangle = \bigotimes_{j=1}^{n} |i_j\rangle$

Such a method introduces problems when considering the operational meaning of the cost function. As discussed above, in the case of the marginal cost functions, the opera-

tional meaning of the Hilbert-Schmidt distance was retained, and simply just focused on reduced states. However, $C_{\mathrm{L}}^{\mathrm{CCPS}}$ no longer follows the definition of the Hilbert-Schmidt distance. By cyclicity of trace, the overlap term may be rewritten as

$$\mathrm{Tr_S}\left[\rho \mathrm{U}_{\tilde{\theta}} \mathrm{H}_{\mathrm{L}}^{\mathrm{CCPS(i)}} \mathrm{U}_{\tilde{\theta}}^{\dagger}\right], \tag{6.26}$$

that is, it is the overlap between the trial state ρ, and the state formed by evolving $H_{\mathrm{L}}^{\mathrm{CCPS}(i)}$ under the unitary $U_{\vec{\theta}}$. Thus, to gain a full understanding of the operational meaning of the cost function, we would need a full understanding of the state obtained (in the most general case) by evolving the local Hamiltonian under $U_{\vec{\theta}}$. Regardless of precise operational meaning, it is still possible to analyse the faithfulness of the cost function, and thus identify cases in which it may prove useful. This cost function is provably faithful in the case of training pure states, with provable bounds on the effectiveness for mixed states, dependent on their purity. In order to prove this, we first state and re-derive a result introduced in [69].

Proposition 6.27. *Let ρ be an arbitrary quantum state, and let $U_{\vec{\theta}}$ be a parametrised unitary (for the preparation of a trial state, σ). For shorthand, denote $\rho_U = U_{\vec{\theta}}^{\dagger} \rho U_{\vec{\theta}}$ and $H_{\mathrm{L}}^{(i)} = H_{\mathrm{L}}^{\mathrm{CCPS}(i)}$. We have that*

$$1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L}}^{(\mathrm{i})}\right] \leq 1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{G}}^{(\mathrm{i})}\right] \leq \mathrm{n}\left(1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L}}^{(\mathrm{i})}\right]\right) \tag{6.28}$$

Proof. We can express the local Hamiltonian in the form $H_{\mathrm{L}}^{(i)} = \frac{1}{n} \sum_j^n H_{\mathrm{L},j}^{(i)}$, where

$$H_{\mathrm{L},j}^{(i)} = |i_j\rangle \langle i_j|_{S_j} \otimes \mathbb{I}_{S_{\bar{j}}} \tag{6.29}$$

are mutually commuting projectors that multiply up to the global Hamiltonian, i.e., $\Pi_{j=1}^n H_{\mathrm{L},j}^{(i)} = H_{\mathrm{G}}^{(i)}$. We can associate events, E_j, with projectors $H_{\mathrm{L},j}$ such that $\Pr[E_j] = \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L,j}}^{(\mathrm{i})}\right]$. Consequently, $\mathrm{Tr}\left[\rho_{\mathrm{U}} \Pi_{\mathrm{j=1}}^{\mathrm{n}} \mathrm{H}_{\mathrm{L,j}}^{(\mathrm{i})}\right] = \Pr\left[\cap_{\mathrm{j=1}}^{\mathrm{n}} \mathrm{E_j}\right]$. Recall that, via the axiomatic definition of probability measures, for any set of events $\mathcal{A}$:]$\{A_1, A_2, ..., A_m\}$, it holds that

$$\Pr\left[\bigcup_{i=1}^{m} A_i\right] \geq \frac{1}{n} \sum_{i=1}^{m} \Pr[A_i]. \tag{6.30}$$

Choosing $A_i = \overline{E_j}$, (where $\overline{E_j}$ denotes the complement of E_j), we see

$$\begin{aligned} \Pr\left[\bigcup_{j=1}^{n} \overline{E_j}\right] &\geq \frac{1}{n}\sum_{j=1}^{n} \Pr\left[\overline{E_j}\right] \\ \implies 1 - \Pr\left[\bigcap_{j=1}^{n} E_j\right] &\geq \frac{1}{n}\sum_{j=1}^{n} \left(1 - \Pr\left[E_j\right]\right) \\ \implies 1 - \Pr\left[\bigcap_{j=1}^{n} E_j\right] &\geq 1 - \frac{1}{n}\sum_{j=1}^{n} \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L,j}}^{(\mathrm{i})}\right] \\ \implies 1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{G}}^{(\mathrm{i})}\right] &\geq 1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L}}^{(\mathrm{i})}\right], \end{aligned} \tag{6.31}$$

which is precisely the LHS of inequality 6.28. In order to prove the remaining inequality, we make use of the union bound, observing that

$$\begin{aligned} \Pr\left[\bigcup_{j=1}^{n} \overline{E_j}\right] &\leq \sum_{j=1}^{n} \Pr\left[\overline{E_j}\right] \\ \implies 1 - \Pr\left[\bigcap_{j=1}^{n} E_j\right] &\leq \sum_{j=1}^{n} \left(1 - \Pr\left[E_j\right]\right) \\ \implies 1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{G}}^{(\mathrm{i})}\right] &\leq n\left(1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L}}^{(\mathrm{i})}\right]\right). \end{aligned} \tag{6.32}$$

Thus, together we have $1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L}}^{(\mathrm{i})}\right] \leq 1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{G}}^{(\mathrm{i})}\right] \leq \mathrm{n}\left(1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L}}^{(\mathrm{i})}\right]\right)$ as required. □

Utilising the above inequality, we can derive a lower bound on $C_{\mathrm{L}}^{\mathrm{CCPS}}$ that guarantees faithfulness for pure states, reliant on the impurity of our target and training states.

Theorem 6.33. *Given two arbitary quantum states, ρ, σ, with σ expressed in the form of the CCPS ansatz, $C_{\mathrm{L}}^{\mathrm{CCPS}}(\rho, \sigma)$ is lower bounded by $C_{\mathrm{G}}(\rho, \sigma)$ in the following way:*

$$n C_{\mathrm{L}}^{\mathrm{CCPS}} \geq C_{\mathrm{G}} - (n-1)(\mathrm{Impurity}(\rho) + \mathrm{Impurity}(\sigma)), \tag{6.34}$$

where $\mathrm{Impurity}(X) := 1 - \mathrm{Tr}\left[\mathrm{X}^2\right]$ *for $X = \rho$ and $X = \sigma$, and $C_{\mathrm{G}} \equiv \|\rho - \sigma\|_2^2$.*

Corollary 6.35. *It follows that if $C_{\mathrm{L}}^{(\mathrm{CCPS})} = 0$, then*

$$(n-1)(\mathrm{Impurity}(\rho) + \mathrm{Impurity}(\sigma) \geq \|\rho - \sigma\|_2^2. \tag{6.36}$$

That is, if both the target and trained states are pure ($\mathrm{Impurity}(\rho) = \mathrm{Impurity}(\sigma) = 0$), then $C_{\mathrm{L}}^{\mathrm{CCPS}} = 0$ implies that $C_{\mathrm{G}} = 0$. More generally, if the impurities of both states are low, then a vanishing local cost function implies that C_{G} is small.

Proof. From (6.28), we have

$$\begin{aligned} & 1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{G}}^{(\mathrm{i})}\right] \leq \mathrm{n}\left(1 - \mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L}}^{(\mathrm{i})}\right]\right) \\ \implies & -\mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{G}}^{(\mathrm{i})}\right] \leq (\mathrm{n}-1) - \mathrm{nTr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L}}^{(\mathrm{i})}\right] \\ \implies & -2\sum_i p_\phi(i)\mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{G}}^{(\mathrm{i})}\right] \leq -2\mathrm{n}\sum_{\mathrm{i}} \mathrm{p}_\phi(\mathrm{i})\mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L}}^{(\mathrm{i})}\right] + 2(\mathrm{n}-1)\sum_{\mathrm{i}} \mathrm{p}_\phi(\mathrm{i}), \end{aligned} \tag{6.37}$$

Adding the purity terms to both sides (and noting that $2(n-1)\sum_i p_\phi(i) = 2(n-1)$) gives

$$\begin{aligned} & C_{\mathrm{G}} \leq \mathrm{Tr}\left[\rho^2\right] + \mathrm{Tr}\left[\sigma^2\right] - 2\mathrm{n}\sum_{\mathrm{i}} \mathrm{p}_\phi(\mathrm{i})\mathrm{Tr}\left[\rho_{\mathrm{U}} \mathrm{H}_{\mathrm{L}}^{(\mathrm{i})}\right] + 2(\mathrm{n}-1) \\ \implies & C_{\mathrm{G}} \leq nC_{\mathrm{L}}^{\mathrm{CCPS}} + (n-1)\left(2 - \mathrm{Tr}\left[\rho^2\right] - \mathrm{Tr}\left[\sigma^2\right]\right) \end{aligned} \tag{6.38}$$

or, equivalently,

$$nC_{\mathrm{L}}^{\mathrm{CCPS}} \geq C_{\mathrm{G}} - (n-1)(\mathrm{Impurity}(\rho) + \mathrm{Impurity}(\sigma). \tag{6.39}$$

Accordingly, if $C_{\mathrm{L}}^{\mathrm{CCPS}} = 0$, then

$$C_{\mathrm{G}} \leq (n-1)(\mathrm{Impurity}(\rho) + \mathrm{Impurity}(\sigma)). \tag{6.40}$$

Therefore, if both the target and training states are pure (states with an impurity of 0), the cost is faithful; i.e., $C_{\mathrm{L}}^{\mathrm{CCPS}} = 0 \implies C_{\mathrm{G}} = 0$. For high purity states, C_{G} remains small, and $C_{\mathrm{L}}^{\mathrm{CCPS}}$ is approximately faithful. □

Note that the construction of $C_{\mathrm{L}}^{\mathrm{CCPS}}$ is equivalent to taking a partitioning of the set of n qubits (into n sets, each with cardinality 1), performing (1-)local measurements on each partition, and taking the average. This idea can be generalised for arbitrary partitionings of the n qubits, allowing for the construction of alternative local cost functions, similarly to 6.19. One simple approach is to perform $\frac{n}{k}$ measurements, on $k < n$ qubits at a time. This of course restricts us to only choosing k such that $k|n$. Defining

$$H_{\mathrm{L}}^{k(i)} := \frac{1}{(n/k)} \sum_{m=1}^{n/k} H_{\mathrm{L},m}^{k(i)}, \tag{6.41}$$

with

$$H_{\mathrm{L},m}^{k(i)} := |i\rangle_{\mathcal{P}_m} \langle i|_{\mathcal{P}_m} \otimes \mathbb{I}_{\overline{\mathcal{P}_m}}, \tag{6.42}$$

where each $\mathcal{P}_m$ contains the indices of the k qubits being measured over by the m-th (k-)local operator, and all $\mathcal{P}_m$ are defined by the choice of partitioning $\mathcal{P}$, such that $\mathcal{P}_1 \cup \mathcal{P}_2 \cup \cdots \cup \mathcal{P}_{n/k}$ spans $\{1, \ldots, n\}$; we can define a k-local cost function specifically for

the CCPS ansatz, as follows:

Definition 6.43. For $k|n$, we have the k-local cost function

$$C_{\mathrm{L}}^{k} = \mathrm{Tr}\left[\rho^2\right] + \sum_{\mathrm{i}} \mathrm{p}_{\phi}(\mathrm{i})^2 - 2\sum_{\mathrm{i}} \mathrm{p}_{\phi}(\mathrm{i})\mathrm{Tr}_{\mathrm{S}}\left[\mathrm{U}_{\tilde{\theta}}^{\dagger}\rho\mathrm{U}_{\tilde{\theta}}\mathrm{H}_{\mathrm{L}}^{\mathrm{k(i)}}\right]. \tag{6.44}$$

Noting that the construction of the k-local cost is similar to that of the 1-local cost (via the use of mutually commuting projectors $H_{\mathrm{L},m}^{k(i)}$), we would expect the k-local cost to be faithful in the same case of the 1-local cost, i.e., for pure states. Indeed, this can be shown, by following the same formula established in proposition 6.27.

Proposition 6.45. *Let ρ be an arbitrary quantum state, and $U_{\vec{\theta}}$ be a parametrised unitary (for the preparation of a trial state, σ). For shorthand, denote $\rho_U = U_{\vec{\theta}}^{\dagger}\rho U_{\vec{\theta}}$. We have that*

$$1 - \mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{L}}^{\mathrm{k(i)}}\right] \leq 1 - \mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{G}}^{\mathrm{(i)}}\right] \leq \frac{\mathrm{n}}{\mathrm{k}}\left(1 - \mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{L}}^{\mathrm{k(i)}}\right]\right) \tag{6.46}$$

Proof. We associate events E_i with the projectors $H_{\mathrm{L},m}^{k(i)}$ such that $\Pr\left[E_i\right] = \mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{L,m}}^{\mathrm{k(i)}}\right]$. Noting that $\mathrm{Tr}\left[\rho_{\mathrm{U}}\Pi_{\mathrm{m=1}}^{\mathrm{n/k}}\mathrm{H}_{\mathrm{L,m}}^{\mathrm{k(i)}}\right] = \Pr\left[\cup_{\mathrm{m=1}}^{\mathrm{n/k}}\mathrm{E_m}\right]$, we have that

$$\begin{aligned}
\Pr\left[\bigcup_{m=1}^{n/k}\overline{E_m}\right] &\geq \frac{1}{(n/k)}\sum_{m=1}^{n/k}\Pr\left[\overline{E_m}\right] \\
\implies 1 - \Pr\left[\bigcap_{m=1}^{n/k}E_m\right] &\geq \frac{1}{(n/k)}\sum_{m=1}^{n/k}\left(1 - \Pr\left[E_m\right]\right) \\
\implies 1 - \mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{G}}^{\mathrm{(i)}}\right] &\geq 1 - \frac{1}{(n/k)}\sum_{m=1}^{n/k}\mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{L,m}}^{\mathrm{k(i)}}\right],
\end{aligned} \tag{6.47}$$

which is equivalent to the LHS of (6.46). We also have that

$$\begin{aligned}
\Pr\left[\bigcup_{m=1}^{n/k}\overline{E_m}\right] &\leq \sum_{m=1}^{n/k}\Pr\left[\overline{E_m}\right] \\
\implies 1 - \Pr\left[\bigcap_{m=1}^{n/k}E_m\right] &\leq \sum_{m=1}^{n/k}\left(1 - \Pr\left[E_m\right]\right) \\
\implies 1 - \mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{G}}^{\mathrm{(i)}}\right] &\leq \frac{n}{k}\left(1 - \mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{L}}^{\mathrm{k(i)}}\right]\right).
\end{aligned} \tag{6.48}$$

Combining the two inequalities results in

$$1 - \mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{L}}^{\mathrm{k(i)}}\right] \leq 1 - \mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{G}}^{\mathrm{(i)}}\right] \leq \frac{\mathrm{n}}{\mathrm{k}}\left(1 - \mathrm{Tr}\left[\rho_{\mathrm{U}}\mathrm{H}_{\mathrm{L}}^{\mathrm{k(i)}}\right]\right). \tag{6.49}$$

□

Comparing (6.46) to (6.28), the newly introduced dependence on k allows us to form a lower bound on the local cost that varies depending on the choice of k, i.e., the choice of "how local" the cost function is.

Theorem 6.50. *Given two arbitrary quantum states, ρ, σ with σ expressed in the form of the CCPS ansatz, $C_{\mathrm{L}}^{k}(\rho, \sigma)$ is lower bounded by $C_{\mathrm{G}}(\rho, \sigma)$ in the following way:*

$$\frac{n}{k} C_{\mathrm{L}}^{k} \geq C_{\mathrm{G}} - \left(\frac{n}{k} - 1\right) \left(\mathrm{Impurity}(\rho) + \mathrm{Impurity}(\sigma)\right). \tag{6.51}$$

Corollary 6.52. *It follows from theorem 6.50 that if $C_{\mathrm{L}}^{k} = 0$, then*

$$\left(\frac{n}{k} - 1\right) \left(\mathrm{Impurity}(\rho) + \mathrm{Impurity}(\sigma)\right) \geq \|\rho - \sigma\|_2^2. \tag{6.53}$$

Proof. The proof is completely analogous to that for theorem 6.33, but with $n \to n/k$. □

Once again, we have that, for states with 0 impurity, a vanishing local cost implies a vanishing global cost. However, the inclusion of the reciprocal of k implies that the k-local cost is "closer to faithful" than the 1-local cost for states with (relatively) low impurity; whilst the cost becomes "increasingly faithful" as k grows. As we should expect, the cost becomes trivially faithful for $k = n$ (equivalently, the global cost). Thus, similarly to the marginal local cost, we could start training on the 1-local cost (or, some convex combination of different k-local costs, with extra weighting given to low k), and increase our dependence on high-k terms as we approach the solution, driving our ansatz towards the global minimum.

Shifting to the framework of the SP ansatz, we can again invoke local costs by replacing the relevant measurement operator with local measurements. First, recall the global cost function as defined for the SP ansatz:

$$\begin{aligned} C_{\mathrm{SP}}(\rho, \sigma) &= \mathrm{Tr}\left[\rho^2\right] + \mathrm{Tr}\left[\sigma^2\right] - 2\mathrm{Tr}\left[\rho\sigma\right] \\ &= \mathrm{Tr}\left[\rho^2\right] + \mathrm{Tr}\left[\sigma\left(\tilde{\theta}, \mathrm{R}\right)^2\right] \\ &\quad - 2d_A \mathrm{Tr}_{\mathrm{SA}}\left[\left(\rho \otimes \frac{\mathbb{I}}{\mathrm{d_A}}\right) \mathrm{U}_{\tilde{\theta}} \left(|0\rangle\langle 0|\right)^{\otimes(\mathrm{n}+\mathrm{n_A})} \mathrm{U}_{\tilde{\theta}}^{\dagger}\right], \end{aligned} \tag{6.54}$$

where we have represented our trial state via its (parametrised) Stinespring representation. Without loss of generality (and for reasons which will become clearer once introducing the local cost equivalent), for analysis we may also express ρ via the Stinespring

representation of its purification, i.e.,

$$\rho = \mathrm{Tr_A}\left[\mathrm{V}_\rho \left(|0\rangle\langle 0|\right)^{\otimes(\mathrm{n}+\mathrm{n_A})} \mathrm{V}_\rho^\dagger\right]. \tag{6.55}$$

As in the CCPS case, we need to identify which terms require a replacement in the (global) measurement operator, $(|0\rangle\langle 0|)^{\otimes(n+n_a)}$, in order to construct a local cost. However, unlike the CCPS case, we calculate the purity term for σ on quantum hardware, as the SP ansatz does not provide a classical description of this term (unlike the CCPS ansatz, where the convex combination $\sum_i p_\phi(i)^2$ completely and classically describes this term). Thus, we need to implement local measurements for (at least) both the overlap term, and the purity of the trial state σ. Noting that, via cyclicity of trace, we have

$$\mathrm{Tr_{SA}}\left[\left(\rho\otimes\frac{\mathbb{I}}{\mathrm{d_A}}\right)\mathrm{U}_{\tilde{\theta}}(|0\rangle\langle 0|)^{\otimes(\mathrm{n}+\mathrm{n_A})}\mathrm{U}_{\tilde{\theta}}^\dagger\right] \equiv \mathrm{Tr_{SA}}\left[\mathrm{U}_{\tilde{\theta}}^\dagger\left(\rho\otimes\frac{\mathbb{I}}{\mathrm{d_A}}\right)\mathrm{U}_{\tilde{\theta}}\left(|0\rangle\langle 0|\right)^{\otimes(\mathrm{n}+\mathrm{n_A})}\right], \tag{6.56}$$

we can indeed treat $(|0\rangle\langle 0|)^{\otimes(n+n_A)}$ as our global measurment operator (despite this being an input ground state for evolution under our training unitary). As such, we propose replacement measurements as follows:

Definition 6.57. By replacing the global measurement with a local operator on the system register, we can define the *Singly-local SP Hamiltonian*,

$$H_{\mathrm{L}}^{\mathrm{S}} = \frac{1}{n}\sum_{j=1}^{n} |0\rangle\langle 0|_{S_j} \otimes \mathbb{I}_{S_{\bar{j}}} \otimes |0\rangle\langle 0|_A. \tag{6.58}$$

Alternatively, we can try to capture an intuitively "more" local set of information, by also making the measurement local on the ancilla register.

Definition 6.59. We define the *Doubly-local SP Hamiltonian* as

$$H_{\mathrm{L}}^{\mathrm{D}} = \frac{1}{n}\sum_{j=1}^{n} |0\rangle\langle 0|_{S_j} \otimes \mathbb{I}_{S_{\bar{j}}} \otimes |0\rangle\langle 0|_{A_j} \otimes \mathbb{I}_{A_{\bar{j}}}. \tag{6.60}$$

To grasp what, if any, operational meaning is to be had in a cost function employing these measurement techniques, before formally defining our local costs, we opt to consider the new form of the overlap term. Without loss of generality, consider the overlap term, $\mathrm{Tr}[\rho\sigma]$, with the global measurement operator replaced with the singly-local SP Hamiltonian,

$$\begin{aligned} &d_A \mathrm{Tr_{SA}}\left[\left(\rho_{\mathrm{S}}\otimes\mathbb{I}/\mathrm{d_A}\right)\mathrm{U}_{\tilde{\theta}}\mathrm{H}_{\mathrm{L}}^{\mathrm{S}}\mathrm{U}_{\tilde{\theta}}^\dagger\right] \\ =&d_A \mathrm{Tr_{SA}}\left[\mathrm{U}_{\tilde{\theta}}^\dagger\left(\rho_{\mathrm{S}}\otimes\mathbb{I}/\mathrm{d_A}\right)\mathrm{U}_{\tilde{\theta}}\mathrm{H}_{\mathrm{L}}^{\mathrm{S}}\right]. \end{aligned} \tag{6.61}$$

By inspection of the first line of (6.61), we see that this term is equivalent to the overlap of $\rho_S \otimes \mathbb{I}/d_A$ and the state obtained by evolving the local Hamiltonian under $U_{\vec{\theta}}$. Clearly, this is not equivalent to the overlap between ρ and σ — the latter is obtained by evolving the all-zero state under $U_{\vec{\theta}}$ (and tracing out the ancilla system), whilst we instead have evolved the maximally mixed state on $n-1$ system qubits along with the all-zero state on one system qubit, and the all-zero state on the ancilla. (For the case of the doubly-local Hamiltonian, this would be the maximally mixed state on $n-1$ and n_A-1 qubits along with the all-zero state on two qubits, one each from the system and ancillary registers). As such, it is unclear what operational meaning remains in the overlap term, being the overlap between our global target state, and some unknown state whose only link to the trial state is that they share a preparation unitary (but differ in initialisations). Whilst this does not necessarily preclude us from using it in training on its own, a combination of this overlap with the (global) purity term on ρ would evidently quickly raise questions in terms of any provable faithfulness: It is not necessarily true that such a cost function would vanish given $\sigma = \rho \iff U_{\vec{\theta}} = V_\rho$, as $V_\rho H_{\mathrm{L}}^{\mathrm{S}} V_\rho^\dagger \neq \rho$. As such, to ensure a vanishing cost when $\sigma = \rho \iff U_{\vec{\theta}} = V_\rho$, we need to also introduce the local measurements into the purity term for ρ.

Definition 6.62. For the SP ansatz, we define a local cost up to choice of singly- versus doubly-local measurement,

$$C_{\mathrm{L}}^{\mathrm{X}} = c_{\mathrm{L}}^{\mathrm{X}}(V_\rho, \rho) + c_{\mathrm{L}}^{\mathrm{X}}(U_{\vec{\theta}}, \sigma) - 2c_{\mathrm{L}}^{\mathrm{X}}(U_{\vec{\theta}}, \rho), \tag{6.63}$$

where we define

$$c_{\mathrm{L}}^{\mathrm{X}}(U_{\vec{\theta}}, \rho) := d_A \mathrm{Tr}_{\mathrm{SA}} \left[\mathrm{U}_{\vec{\theta}}^\dagger (\rho_{\mathrm{S}} \otimes \mathbb{I}/\mathrm{d}_{\mathrm{A}}) \mathrm{U}_{\vec{\theta}} \mathrm{H}_{\mathrm{L}}^{\mathrm{X}} \right] \tag{6.64}$$

with a choice of $X = S(D)$ producing the *singly-* (*doubly-*)*local cost function*.

We note that each term can be efficiently calculated on quantum hardware via the use of Loschmidt-Echo style circuits. However, the first term, $c_{\mathrm{L}}^{\mathrm{X}}(V_\rho, \rho)$ will not be measurable in practicality, as we do not have access to the purifying unitary V_ρ (in fact, this is exactly what we are trying to learn via the SP ansatz approach). Yet this does not preclude us from using this cost function, as this term remains constant throughout and does not contribute to the gradient change of the cost function, allowing it to be neglected with no effect on the optimisation procedure. Noting that we once again have a Hamiltonian comprised of mutually commuting projections, it is natural to wonder if this cost function is provably faithful for pure states, as the CCPS-ansatz based local cost function was. However, the presence of the d_A factor in our calculations causes such analysis to fail. Despite this, it is still possible to prove faithfulness for tensor-product states in the SP picture.

Theorem 6.65. *$C_{\mathrm{L}}^{\mathrm{D}}$ is faithful to the global Hilbert-Schmidt cost for unentangled (tensor-product) states.*

Proof. Given an unentangled target state, we can express it in the form $\rho = \bigotimes_{j=1}^{n} \rho_{S_j}$. The preparation unitary for our trial state σ can, w.l.o.g., take the form $U_{\vec{\theta}} = \bigotimes_{j=1}^{n} U_{S_j A_j}$, i.e., the preparation of each qubit is independent from the preparation of all other qubits (as would be expected for an unentangled state). Thus, the overlap can be written as

$$\begin{aligned}
c_{\mathrm{L}}^{\mathrm{D}}(U_{\vec{\theta}}, \rho) &= d_A \mathrm{Tr}_{\mathrm{SA}} \left[\left(\bigotimes_{\mathrm{k}=1}^{\mathrm{n}} \mathrm{U}_{\mathrm{S_k A_k}}^{\dagger} \left(\bigotimes_{\mathrm{k}=1}^{\mathrm{n}} \rho_{\mathrm{S_k}} \otimes \frac{\mathbb{I}_{\mathrm{A}}}{\mathrm{d_A}} \right) \bigotimes_{\mathrm{k}=1}^{\mathrm{n}} \mathrm{U}_{\mathrm{S_k A_k}} \right) \mathrm{H}_{\mathrm{L}}^{\mathrm{D}} \right] \\
&= \frac{d_A}{n} \sum_{j=1}^{n} \mathrm{Tr}_{\mathrm{SA}} \times \\
&\left[\left(\bigotimes_{k=1}^{n} U_{S_k A_k}^{\dagger} \left(\bigotimes_{k=1}^{n} \rho_{S_k} \otimes \frac{\mathbb{I}_A}{d_A} \right) \bigotimes_{k=1}^{n} U_{S_k A_k} \right) \left(|0\rangle\langle 0|_{S_j} \otimes \mathbb{I}_{S_{\bar{j}}} \otimes |0\rangle\langle 0|_{A_j} \otimes \mathbb{I}_{A_{\bar{j}}} \right) \right] \\
&= \frac{1}{n} \sum_{j=1}^{n} \mathrm{Tr}_{\mathrm{S_j A_j}} \left[\left(\mathrm{U}_{\mathrm{S_j A_j}}^{\dagger} \left(\rho_{\mathrm{S_j}} \otimes \mathbb{I}_{\mathrm{A_j}} \right) \mathrm{U}_{\mathrm{S_j A_j}} \right) \left(|0\rangle\langle 0|_{\mathrm{S_j A_j}} \right) \right] \times \\
&\mathrm{Tr}_{\mathrm{S_{\bar{j}} A_{\bar{j}}}} \left[\left(\bigotimes_{\mathrm{k} \neq \mathrm{j}} \rho_{\mathrm{S_k}} \otimes \mathbb{I}_{\mathrm{A_{\bar{j}}}} \right) \right] \\
&= \frac{d_A}{2n} \sum_{j=1}^{n} \mathrm{Tr}_{\mathrm{S_j A_j}} \left[\mathrm{U}_{\mathrm{S_j A_j}}^{\dagger} \left(\rho_{\mathrm{S_j}} \otimes \mathbb{I}_{\mathrm{A_j}} \right) \mathrm{U}_{\mathrm{S_j A_j}} |0\rangle\langle 0|_{\mathrm{S_j A_j}} \right] \\
&= \frac{d_A}{2n} \sum_{j=1}^{n} \mathrm{Tr} \left[\rho_{\mathrm{S_j}} \sigma_{\mathrm{S_j}} \right].
\end{aligned} \tag{6.66}$$

Thus, for tensor product states, we have that

$$C_{\mathrm{L}}^{\mathrm{D}} = \frac{d_A}{2n} \left(\sum_{j=1}^{n} \mathrm{Tr} \left[\rho_{\mathrm{j}} \rho_{\mathrm{j}} \right] + \sum_{\mathrm{j}=1}^{\mathrm{n}} \mathrm{Tr} \left[\sigma_{\mathrm{j}} \sigma_{\mathrm{j}} \right] - 2 \sum_{\mathrm{j}=1}^{\mathrm{n}} \mathrm{Tr} \left[\rho_{\mathrm{j}} \sigma_{\mathrm{j}} \right] \right). \tag{6.67}$$

Noting that

$$\sum_{j=1}^{n} \mathrm{Tr} \left[\rho_{\mathrm{j}} \rho_{\mathrm{j}} \right] + \sum_{\mathrm{j}=1}^{\mathrm{n}} \mathrm{Tr} \left[\sigma_{\mathrm{j}} \sigma_{\mathrm{j}} \right] - 2 \sum_{\mathrm{j}=1}^{\mathrm{n}} \mathrm{Tr} \left[\rho_{\mathrm{j}} \sigma_{\mathrm{j}} \right] = \sum_{\mathrm{j}=1}^{\mathrm{n}} |\rho_{\mathrm{j}} - \sigma_{\mathrm{j}}|_2^2, \tag{6.68}$$

we have that

$$C_{\mathrm{L}}^{\mathrm{D}} \propto \sum_{j=1}^{n} |\rho_j - \sigma_j|_2^2, \tag{6.69}$$

where each term of the summation is greater than or equal to zero, with equality if and only if $\rho_j = \sigma_j$, therefore the cost may only vanish if and only if $\rho_j = \sigma_j$ for all j. □

It is worth noting that the inability to construct a local variant of the cost function for the SP ansatz is not due to the inherent mathematical characteristics of the singly- and doubly-local cost functions presented above, but rather due to the practical limitations when implementing the algorithm. That we do not have access to the purification of ρ (as it is in fact this precise purification that we are hoping to learn) precludes faithfulness. A faithful (but impractical for implementation) cost function may be formulated by introducing (either) local Hamiltonian into each state referenced in the cost function, i.e.,

Definition 6.70. Impractical local cost.

$$C_{\mathrm{Im}}^{\mathrm{X}} = c_{\mathrm{L}}^{\mathrm{X}}\left(V_\rho, V_\rho\right) + c_{\mathrm{L}}^{\mathrm{X}}\left(U_{\vec{\theta}}, U_{\vec{\theta}}\right) - 2c_{\mathrm{L}}^{\mathrm{X}}\left(U_{\vec{\theta}}, V_\rho\right). \tag{6.71}$$

Here, by momentarily ignoring the practical aspects of training, we have the easily defineable operational meaning of the Hilbert Schmidt distance between the state found via the purifications V_ρ and $U_{\vec{\theta}}$, leading to trivial faithfulness.

Thus, the existence of a local formulation of the Hilbert-Schmidt distance that is suitable for training on current and near-term quantum hardware (in tandem with classical counterparts) remains an open question. Nonetheless, an algorithm that has great potential for success with current hardware is proposed. Analysis of performance of the algorithm on quantum hardware may be found in [39].

Chapter 7

Discussion and Future Work

In this work, the use of manipulating and interpreting quantum sources of information is explored in two ways. The principally experimental chapters (4, 5) explore the interpretation of information which, despite being classical, has a high degree of randomness stemming from quantum materials. It is demonstrated that semiconducting CuInS/ZnS core/shell colloidal quantum dot based ink is a viable candidate for PUF devices aimed towards ubiquitous use. Of the two algorithms investigated in this work, Reduced Modified Local Binary Patterns (R-MLBP) and Adapted High Boost (AHB), R-MLBP is judged to be more suitable for use in workflows employing smartphones in general settings, primarily due to the easily separable intra- and inter-Hamming distance distributions obtained for fingerprints generated from images of devices in various settings. R-MLBP's built in noise resistance allows for confidence in a selection of parameters for wide use, with responses generated in different environments bearing more similarity under R-MLBP.

A novel type of optical PUF is explored, utilising two layers of quanutm dot-based ink, each with a different peak emission wavelength. Alongside this, a "one-to-many" challenge-response mechanism classification is proposed for this type of PUF, and others explored in the future. It's assessed that both inks are capable of producing highly unique and repeatable fingerprints when compared to other devices, and that comparison of fingerprints from different isolated emissions of the same device reveals a high degree of uniqueness, with some guaranteed low level of correlation. When comparing the two algorithms, AHB appears to capture slightly more correlated fingerprints.

For both ubiquitous PUF designs as analysed in chapter 4, and One-to-Many systems as analysed in chapter 5, R-MLBP is determined to be more suited to the task of interrogating quantum dot based OPUFs. In general, AHB appears to suffer from a poor fingerprinting binary bias, making it unsuitable for the preparation of crpytographic keys as is, although tweaks to the algorithm and/or pre-/post-processing of the image/fingerprint may remedy this. In the ubiquitous case, R-MLBP tends to produce

more easily-separable intra- and inter-Hamming distributions, allowing for greater confidence of no false-matching in wide use. Further, R-MLBP also benefits from a greater consistency in performance trends across different PUF devices, making it a strong candidate for future use and further research in PUF identification.

Chapter 6 explores how we can learn information about a given quantum system, by directly making use of quantum theory. The work principally focuses on the mathematical analysis and treatment of cost functions used in variational quantum algorithms, particularly ones utilising the Hilbert-Schmidt distance as an operationally meaningful metric. Barren plateaus, a phenomena present in cost function landscapes for large-scale problems run on deep hardware, plague current-term research into VQAs. As an attempt to mitigate BPs, this work sought to find a cost function in which, via either the introduction of local states or replacement of global operators in the HS distance with local alternaives, the Hilbert-Schmidt distance was generalised into a local equivalent. Whilst proposed cost functions were seperably proved faithful for the (i) pure state and (ii) tensor product state cases, no cost function could be found that was suitable and faithful for the more general highly entangled and highly mixed state case(s).

7.1 Future Work

7.1.1 Authentication

For widespread use, further investigation into the effects of different uncontrollable environmental factors is needed. Principally, the exploring of the change of angle of incidence (and capture), as well as the introduction of unsteadiness present in human handling of the phone and the PUF device itself is of interest (and is discussed briefly in appendix A). Whilst R-MLBP demonstrated a strong ability to reliably extract unique fingerprints in a variety of lighting conditions, it is necessary to collect more data, involving ambient lighting different warmths, hues, and stability. It is expected that the introduction of an automated calibration system for the choosing of target capture settings such as ISO, exposure, and aperture would allow for an expansion of useable environments for the PUF devices, and an accompanying algorithm. Within this work, characterisation of the sensor used did not take place. Whilst in the pursuit of ubiquitous use, such properties will be assumed unknown (due to the need to operate on devices which may have different sensors, and which may suffer different defects), future work should consider such characterisations, allowing for an understanding of their effect (and extents of effect), such that mitigations can be deployed to increase universality. Whilst the work presented in this thesis extends a weak PUF to a "One-to-X" PUF, an investigation into the inde-

pendence of responses garnered from changing the incident peak wavelength of excitation (and thus, presenting a different challenge) could allow for the creation of "One-to-X" PUFs within the strong PUF domain.

For the HPUFs examined, their lack of close correspondence in fingerprints for devices that appear visually similar is a potential interesting avenue for future research. Understanding the connection between micro-scale features, and how they may or may not be tied into specifics of variable lacquer distribution may allow for further control on the level of correlation present in isolated fingerprints of such devices, allowing fine tuning of devices for specific protocols. With regards to the impact of lacquer disposition on correlation, experiments investigating PUF response behaviour in the presence of ambient light, and/or without a filter on the incident light may be of interest, to explore how reflected light may impact the behaviour of the conjectured thinner areas of lacquer on devices. To this end, it would also be of interest to replicate the experiment on a smartphone camera, adjoining this work with the work towards ubiquitous PUFs, by potentially isolating emission based on single channels of the captured RGB image. In order to investigate, and potentially eliminate, this conjectured lacquer phenomenon, it would be of interest to examine how variation of fabrication method impacts both correlation of fingerprints, and appearance of the device under a microscope. The use of a draw-down applicator for the ink would eliminate the uncontrollable human element introduced in the handstamp method used for this work. Future studies on the clustering of quantum dots in lacquer, particularly with a focus on how fabrication methods may influence the clustering would be of interest, for potential impact on correlation.

Another method of influencing output correlation in HPUF devices is the changing of peak emission of each dot type. Peak emissions may be brought closer to increase likelihood of correlation due to overlaps in emission spectra, or be pushed further apart in an effort to create completely separable fingerprints. Additionally, multiple dots can be used, to create a PUF utilising One-to-X challenge-response mechanism, where X is the number of different quantum dot inks.

Further, this work only considers dual emission quantum dot OPUFs fabricated by the deposition of two different inks onto a surface. Whilst this is believed to be linked to the appearance of correlations in output fingers, as discussed, the analysis of devices fabricated by mixing a single ink comprising of two types of quantum dots of differing sizes would be of special interest, allowing for a comparison of output correlation to the case examined here, and of general interest for the development of new optical PUFs.

More generally, further work into additional layers of security that can be provided would be of interest. Whilst the work in [44] is discussed in this thesis as potential secondary check, it would be interesting to see if the non-linearity that is observable may

be directly paired with the PUF's unique output fingerprint. That such secondary checks exist for non-linearity suggests it is also worth trying to exploit other optical properties of materials more generally, for instance, introducing a phosphorescent layer with a long-time decay could allow for the introduction of a temporal element, as well as a spatial element, to optical PUF-based authentication schemes.

Both types of PUF examined in this thesis utilise small sample sizes, in comparison to potential real-world use cases. Future work would benefit from a larger number of samples, allowing for reaffirmation of the statistics reported in this thesis. Such work would also aid in formalising standards for assessing intra-and inter-hamming distributions, and their related metrics, allowing for a comparison of theoretical false rates against experimental false rates.

7.1.2 Compilation

Quantum compilation as a task is not limited to quantum states only. By learning how to compile any arbitrary quantum circuit, it is then possible to compile any arbitrary quantum processes. One of particular interest, and a natural extension of state compilation, is *quantum channel compilation*. Given a compilation algorithm that can accurately learn any arbitrary mixed state, one can accurately learn any arbitrary quantum channel by exploiting the *Choi-Jamiołkowski isomorphism*, which shows that any quantum channel, equivalent to a completely positive map, has a dual object taking the form of a density matrix, which may be interpreted as a quantum state. In the discussion of channel-state duality, the state associated with a given channel is often referred to as the Choi state. Learning a channel's Choi state is then equivalent to learning the channel itself, following some post-processing. Further, given a PUF whose output is a quantum state, a successful compilation algorithm would provide a method of learning the output, allowing for comparison against the registered (expected) output for the challenge and device.

An open question from this work is whether or not a faithful and practical local equivalent of the Hilbert-Schmidt cost exists for hybrid-based training. Whilst one could not be provably found, any efforts to provide a no-go theorem failed. To this end, either a proof of the existence of a satisfactory cost function, or a no-go theorem for its existence based on properties of subsystems of its elementary operators would be useful for continuing research in efficient and trainable state compilers. Further, this question may have links to the field of quantum foundations, should the reasoning for a no-go theorem be based on operational aspects of mixed states. If it is known that mixed states can not be compiled with guarantee, this sheds some light on what can be learned about quantum objects overall.

Appendix A

Towards Implementation

In chapter 4, the use of smartphones to interrogate (and record the subsequent optical response of) quantum dot-based PUF devices is demonstrated, under tightly controlled conditions. Whilst the initial controls on ambient light are relaxed, the issue of human handling, of both device and tags, is not investigated. This appendix provides a brief, preliminary look at how issues of varying angle and distance between the smartphone and PUF device can be mitigated.

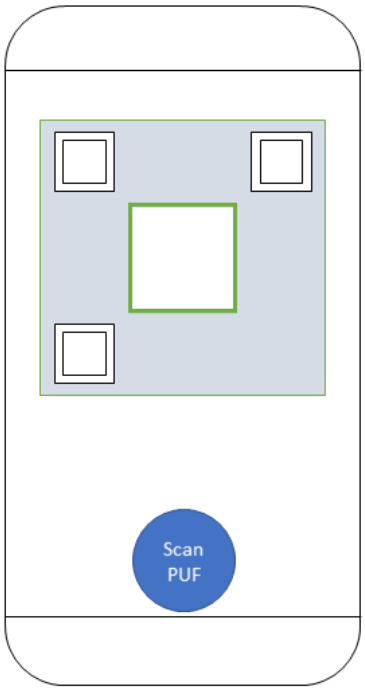

Figure A.1: Simple schematic to prompt a smartphone user to interrogate a PUF device. The user is prompted to align the QR's eyes with the three corner squares, such that captured images are taken at distances and angles that are consistent across multiple challenges.

Given that a custom-designed app would need to be used in order to both trigger the capturing of a CRP output, and the execution of an authentication protocol based on it, one solution is to control the user's capturing behaviour where possible, via instructions within the app. These instructions can take the form of a simple text dialog, or be combined with a visual prompt, i.e., a custom viewfinder shown on screen (a simple example of which is provided in figure A.1), exploiting the landmarks of the QR (the eyes used as part of the QR detection process) to ensure that the user aligns the tag consistently, to produce an image of expected perspective, resolution, and orientation.

However, some human error is likely to persist, resulting in variations of angle, that may potentially lead to warping, particularly around the edges of the quantum dot patch. To this end, the eyes of the QR may again be used, as anchor points for a standard perspective correction algorithm, restoring a "face-on" angle to the imaged PUFs, as demonstrated in figure A.2

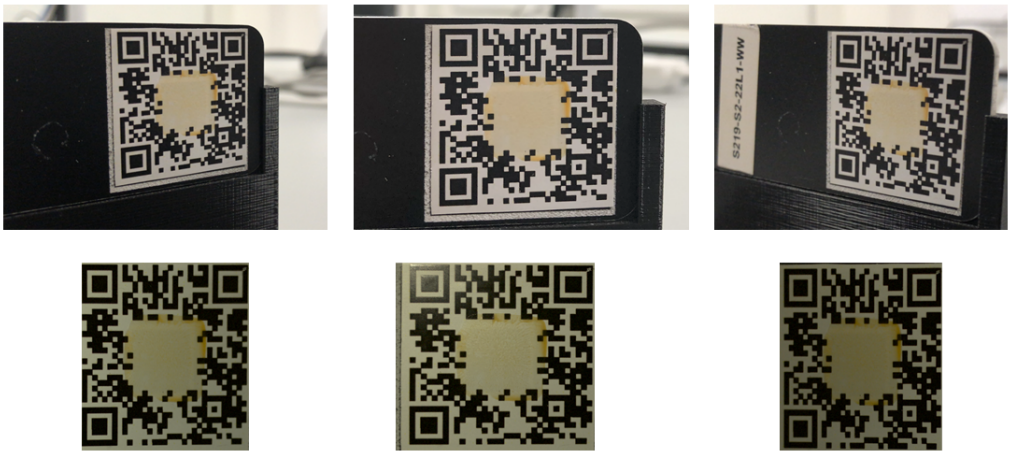

Figure A.2: Example inputs (top row) and outputs (bottom row) for a custom designed perspective correction algorithm, utilising the QR eyes as anchor points to produce more consistent input images for fingerprinting.

Three images were taken at each angle, and the quantum dot patch was manually cropped in each, before resizing down to 175×175 pixels and processing through R-MLBP via the steps defined in 3. From these R-MLBP fingerprint outputs, an average Hamming distance of 0.31 was calculated, suggesting that even without a visual prompt to minimise human error, relatively reliable fingerprints are obtainable. However, this is a significant increase compared to the results presented in chapter 4 —suggesting the need for further work to investigate this, and test if further image correction (e.g., by homographical transformations) can help mitigate this increase.

Bibliography

[1] *Optical Spectra of CuInS/ZnS Nanocrystals Coated with Oleylamine Ligands*, 1 2022.

[2] Scott Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing*, pages 325–338, 2018.

[3] Samet Akcay and Toby Breckon. Towards automatic threat detection: A survey of advances of deep learning within x-ray security imaging. *Pattern Recognition*, 122:108245, 2022.

[4] Martin R Albrecht, Miloš Prokop, Yixin Shen, and Petros Wallden. Variational quantum solutions to the shortest vector problem. *Quantum*, 7:933, 2023.

[5] Muhammad Naveed Aman, Uzair Javaid, and Biplab Sikdar. A privacy-preserving and scalable authentication protocol for the internet of vehicles. *IEEE Internet of Things Journal*, 8(2):1123–1139, 2020.

[6] Andrew Arrasmith, Zoë Holmes, Marco Cerezo, and Patrick J Coles. Equivalence of quantum barren plateaus to cost concentration and narrow gorges. *Quantum Science and Technology*, 7(4):045015, 2022.

[7] Gaurang Bansal, Naren Naren, Vinay Chamola, Biplab Sikdar, Neeraj Kumar, and Mohsen Guizani. Lightweight mutual authentication protocol for V2G using physical unclonable function. *IEEE Transactions on Vehicular Technology*, 69(7):7234–7246, 2020.

[8] Adriano Barenco, Charles H Bennett, Richard Cleve, David P DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Physical review A*, 52(5):3457, 1995.

[9] DW Bauder. An anti-counterfeiting concept for currency systems. *Sandia National Labs, Albuquerque, NM, Tech. Rep. PTK-11990*, 1983.

[10] Bryce E Bayer. Color imaging array, U.S. Patent 3 971 065, July 1976.

[11] Richard Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966.

[12] Marcello Benedetti, Mattia Fiorentini, and Michael Lubasch. Hardware-efficient variational quantum algorithms for time evolution. *Physical Review Research*, 3(3):033083, 2021.

[13] E Bernstein and U Vazirani. Quantum computation complexity. *SIAM J. Comput*, 26:1411, 1997.

[14] Andrew Blance and Michael Spannowsky. Quantum machine learning for particle physics using a variational quantum classifier. *Journal of High Energy Physics*, 2021(2):1–20, 2021.

[15] Kyle Booth, Minh Do, J Beck, Eleanor Rieffel, Davide Venturelli, and Jeremy Frank. Comparing and integrating constraint programming and temporal planning for quantum circuit compilation. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 28, pages 366–374, 2018.

[16] MD Bowdrey, JA Jones, Emanuel Knill, and R Laflamme. Compiling gate networks on an ising quantum computer. *Physical Review A*, 72(3):032315, 2005.

[17] G Brassard and Charles H Bennett. Quantum cryptography: Public key distribution and coin tossing. In *International conference on computers, systems and signal processing*, pages 175–179, 1984.

[18] Yuan Cao, Wanyi Liu, Lan Qin, Bingqiang Liu, Shuai Chen, Jing Ye, Xianzhao Xia, and Chao Wang. Entropy sources based on silicon chips: True random number generator and physical unclonable function. *Entropy*, 24(11):1566, 2022.

[19] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature communications*, 12(1):1–12, 2021.

[20] Isaac L Chuang, Neil Gershenfeld, and Mark Kubinec. Experimental implementation of fast quantum searching. *Physical review letters*, 80(15):3408, 1998.

[21] Isaac L Chuang, Lieven MK Vandersypen, Xinlan Zhou, Debbie W Leung, and Seth Lloyd. Experimental realization of a quantum algorithm. *Nature*, 393(6681):143–146, 1998.

[22] Lukasz Cincio, Yiğit Subaşı, Andrew T Sornborger, and Patrick J Coles. Learning the quantum algorithm for state overlap. *New Journal of Physics*, 20(11):113022, 2018.

[23] David G Cory, Amr F Fahmy, and Timothy F Havel. Ensemble quantum computing by NMR spectroscopy. *Proceedings of the National Academy of Sciences*, 94(5):1634–1639, 1997.

[24] Marcus Cramer, Martin B Plenio, Steven T Flammia, Rolando Somma, David Gross, Stephen D Bartlett, Olivier Landon-Cardinal, David Poulin, and Yi-Kai Liu. Efficient quantum state tomography. *Nature communications*, 1(1):149, 2010.

[25] G Mauro D'Ariano, Matteo GA Paris, and Massimiliano F Sacchi. Quantum tomography. *Advances in imaging and electron physics*, 128:206–309, 2003.

[26] John Daugman. How iris recognition works. In *The essential guide to image processing*, pages 715–739. Elsevier, 2009.

[27] John G Daugman. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *JOSA A*, 2(7):1160–1169, 1985.

[28] Francesca Dell'Acqua. Lux et vitrum": the evolution of stained glass from the late Roman Empire to the Gothic Age. In Leo S. Olschki, editor, *When glass matters : studies in the history of science of art from Graeco-Roman antiquity to Early Modern Era*. Metropolitan Museum of Art, 2004.

[29] Jeroen Delvaux, Dawu Gu, Ingrid Verbauwhede, Matthias Hiller, and Meng-Day Yu. Efficient fuzzy extraction of PUF-induced secrets: Theory and applications. In *International Conference on Cryptographic Hardware and Embedded Systems*, pages 412–431. Springer, 2016.

[30] David Deutsch. Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985.

[31] David Deutsch. Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985.

[32] David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 439(1907):553–558, 1992.

[33] David Elieser Deutsch, Adriano Barenco, and Artur Ekert. Universality in quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 449(1937):669–677, 1995.

[34] Yevgeniy Dodis, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *Advances in Cryptology-EUROCRYPT 2004: International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004. Proceedings 23*, pages 523–540. Springer, 2004.

[35] Mila Efimenko, Alexander Ignatev, and Konstantin Koshechkin. Review of medical image recognition technologies to detect melanomas using neural networks. *BMC bioinformatics*, 21(11):1–7, 2020.

[36] Alexander L Efros and Louis E Brus. Nanocrystal quantum dots: from discovery to modern development. *ACS nano*, 15(4):6192–6210, 2021.

[37] Artur K Ekert. Quantum cryptography based on bell's theorem. *Physical review letters*, 67(6):661, 1991.

[38] Suguru Endo, Jinzhao Sun, Ying Li, Simon C Benjamin, and Xiao Yuan. Variational quantum simulation of general processes. *Physical Review Letters*, 125(1):010501, 2020.

[39] Nic Ezzell, Elliott M Ball, Aliza U Siddiqui, Mark M Wilde, Andrew T Sornborger, Patrick J Coles, and Zoë Holmes. Quantum mixed state compiling. *arXiv preprint arXiv:2209.00528*, 2022.

[40] Nic Ezzell, Zoë Holmes, and Patrick J Coles. The quantum low-rank approximation problem. *arXiv preprint arXiv:2203.00811*, 2022.

[41] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.

[42] Dmitry A Fedorov, Bo Peng, Niranjan Govind, and Yuri Alexeev. VQE method: a short survey and recent developments. *Materials Theory*, 6(1):1–21, 2022.

[43] Richard P Feynman et al. Simulating physics with computers. *Int. j. Theor. phys*, 21(6/7), 2018.

[44] Matthew J Fong, Christopher S Woodhead, Nema M Abdelazim, Daniel C Abreu, Angelo Lamantia, Elliott M Ball, Kieran Longmate, David Howarth, Benjamin J

Robinson, Phillip Speed, et al. Using intrinsic properties of quantum dots to provide additional security when uniquely identifying devices. *Scientific Reports*, 12(1):16919, 2022.

[45] Dennis Gabor. Theory of communication. part 1: The analysis of information. *Journal of the Institution of Electrical Engineers-part III: radio and communication engineering*, 93(26):429–441, 1946.

[46] Zaixin Gan, Feiliang Chen, Qian Li, Mo Li, Jian Zhang, Xueguang Lu, Lu Tang, Zhao Wang, Qiwu Shi, Weili Zhang, et al. Reconfigurable optical physical unclonable functions enabled by VO2 nanocrystal films. *ACS Applied Materials & Interfaces*, 14(4):5785–5796, 2022.

[47] Blaise Gassend, Dwaine Clarke, Marten Van Dijk, and Srinivas Devadas. Silicon physical random functions. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pages 148–160, 2002.

[48] Neil A Gershenfeld and Isaac L Chuang. Bulk spin-resonance quantum computation. *science*, 275(5298):350–356, 1997.

[49] James Power Gordon. Quantum effects in communications systems. *Proceedings of the IRE*, 50(9):1898–1908, 1962.

[50] Daniel Gottesman and Hoi-Kwong Lo. From quantum cheating to quantum security. *Physics Today*, 53(11):22–27, 2000.

[51] Arseni Goussev, Rodolfo A Jalabert, Horacio M Pastawski, and Diego Wisniacki. Loschmidt echo. *arXiv preprint arXiv:1206.6348*, 2012.

[52] Brian C Grubel, Bryan T Bosworth, Michael R Kossey, Hongcheng Sun, A Brinton Cooper, Mark A Foster, and Amy C Foster. Silicon photonic physical unclonable function. *Optics Express*, 25(11):12710–12721, 2017.

[53] Stuart Hadfield, Zhihui Wang, Bryan O'gorman, Eleanor G Rieffel, Davide Venturelli, and Rupak Biswas. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, 12(2):34, 2019.

[54] Ghaith Hammouri, Aykutlu Dana, and Berk Sunar. CDs have fingerprints too. In *CHES*, volume 9, pages 348–362. Springer, 2009.

[55] Aram W Harrow, Benjamin Recht, and Isaac L Chuang. Efficient discrete approximations of quantum gates. *Journal of Mathematical Physics*, 43(9):4445–4451, 2002.

[56] Ralph VL Hartley. Transmission of information 1. *Bell System technical journal*, 7(3):535–563, 1928.

[57] Jane Hayward. *English and French medieval stained glass in the collection of the Metropolitan Museum of Art*, volume 1. Metropolitan Museum of Art, 2003.

[58] Dong-Chen He and Li Wang. Texture unit, texture spectrum, and texture analysis. *IEEE transactions on Geoscience and Remote Sensing*, 28(4):509–512, 1990.

[59] Ryan Helinski, Dhruva Acharyya, and Jim Plusquellic. A physical unclonable function defined using power distribution system equivalent resistance variations. In *Proceedings of the 46th Annual Design Automation Conference*, pages 676–681, 2009.

[60] Zhaoying Hu, Jose Miguel M Lobez Comeras, Hongsik Park, Jianshi Tang, Ali Afzali, George S Tulevski, James B Hannon, Michael Liehr, and Shu-Jen Han. Physically unclonable cryptographic primitives using self-assembled carbon nanotubes. *Nature nanotechnology*, 11(6):559–565, 2016.

[61] K Igeta and Y Yamamoto. Quantum mechanical computers with single atom and photon fields. In *International Quantum Electronics Conference*, page TuI4. Optica Publishing Group, 1988.

[62] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography. In *30th Annual Symposium on Foundations of Computer Science*, pages 230–235. IEEE Computer Society, 1989.

[63] O Ivanova, A Elliott, T Campbell, and CB Williams. Unclonable security features for additive manufacturing. *Additive Manufacturing*, 1:24–31, 2014.

[64] H Shelton Jacinto, A Matthew Smith, and Nader I Rafla. Utilizing a fully optical and reconfigurable PUF as a quantum authentication mechanism. *OSA Continuum*, 4(2):739–747, 2021.

[65] Jonathan A Jones and Michele Mosca. Implementation of a quantum algorithm on a nuclear magnetic resonance quantum computer. *The Journal of chemical physics*, 109(5):1648–1653, 1998.

[66] Tyson Jones and Simon C Benjamin. Quantum compilation and circuit optimisation via energy dissipation. *arXiv preprint arXiv:1811.03147*, 2018.

[67] Hyunho Kang, Yohei Hori, Toshihiro Katashita, Manabu Hagiwara, and Keiichi Iwamura. Cryptographie key generation from PUF data using efficient fuzzy extractors. In *16th International conference on advanced communication technology*, pages 23–26. IEEE, 2014.

[68] Justin Ker, Lipo Wang, Jai Rao, and Tchoyoson Lim. Deep learning applications in medical image analysis. *Ieee Access*, 6:9375–9389, 2017.

[69] Sumeet Khatri, Ryan LaRose, Alexander Poremba, Lukasz Cincio, Andrew T Sornborger, and Patrick J Coles. Quantum-assisted quantum compiling. *Quantum*, 3:140, 2019.

[70] Min Seong Kim and Gil Ju Lee. Visually hidden, self-assembled porous polymers for optical physically unclonable functions. *ACS Applied Materials & Interfaces*, 2023.

[71] Rolf Landauer. Information is physical. *Physics Today*, 44(5):23–29, 1991.

[72] Margaret Larimer. *Experiencing Divine Light: The Stained Glass Windows of the Miracles of the Virgin at Orsanmichele, Florence*. University of California, Davis, 2014.

[73] Jae W Lee, Daihyun Lim, Blaise Gassend, G Edward Suh, Marten Van Dijk, and Srinivas Devadas. A technique to build a secret key in integrated circuits for identification and authentication applications. In *2004 Symposium on VLSI Circuits. Digest of Technical Papers (IEEE Cat. No. 04CH37525)*, pages 176–179. IEEE, 2004.

[74] Mingle Liao, Jun Yuan, Feng Huang, Pidong Wang, Wenjie Wang, Siyuan Luo, and Yao Yao. On-chip silicon optical scattering physical unclonable function towards hardware security. *Journal of Lightwave Technology*, 2022.

[75] Yang Liu, Fei Han, Fushan Li, Yan Zhao, Maosheng Chen, Zhongwei Xu, Xin Zheng, Hailong Hu, Jianmin Yao, Tailiang Guo, et al. Inkjet-printed unclonable quantum dot fluorescent anti-counterfeiting labels with artificial intelligence authentication. *Nature communications*, 10(1):2409, 2019.

[76] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.

[77] Shih-Chung B Lo, Heang-Ping Chan, Jyh-Shyan Lin, Huai Li, Matthew T Freedman, and Seong K Mun. Artificial convolution neural network for medical image pattern recognition. *Neural networks*, 8(7-8):1201–1214, 1995.

[78] Kieran D Longmate, Nema M Abdelazim, Elliott M Ball, Joonas Majaniemi, and Robert J Young. Improving the longevity of optically-read quantum dot physical unclonable functions. *Scientific Reports*, 11(1):10999, 2021.

[79] Roel Maes, Anthony Van Herrewege, and Ingrid Verbauwhede. PUFKY: A fully functional PUF-based cryptographic key generator. In *Cryptographic Hardware and Embedded Systems–CHES 2012: 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings 14*, pages 302–319. Springer, 2012.

[80] Abhranil Maiti, Vikash Gunreddy, and Patrick Schaumont. A systematic method to evaluate and compare the performance of physical unclonable functions. *Embedded systems design with FPGAs*, pages 245–267, 2013.

[81] Dmitri Maslov, Gerhard W Dueck, D Michael Miller, and Camille Negrevergne. Quantum circuit simplification and level compaction. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 27(3):436–444, 2008.

[82] Sam McArdle, Tyson Jones, Suguru Endo, Ying Li, Simon C Benjamin, and Xiao Yuan. Variational ansatz-based quantum simulation of imaginary time evolution. *npj Quantum Information*, 5(1):75, 2019.

[83] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature communications*, 9(1):4812, 2018.

[84] Thomas McGrath, Ibrahim E Bagci, Zhiming M Wang, Utz Roedig, and Robert J Young. A PUF taxonomy. *Applied Physics Reviews*, 6(1):011303, 2019.

[85] Jae-Won Nam, Ju-Hyeok Ahn, and Jong-Phil Hong. Compact SRAM-Based PUF chip employing body voltage control technique. *IEEE Access*, 10:22311–22319, 2022.

[86] Richard Yew Fatt Ng, Yong Haur Tay, and Kai Ming Mok. A review of iris recognition algorithms. In *2008 International Symposium on Information Technology*, volume 2, pages 1–7. IEEE, 2008.

[87] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

[88] Harry Nyquist. Certain factors affecting telegraph speed. *Transactions of the American Institute of Electrical Engineers*, 43:412–422, 1924.

[89] Timo Ojala, Matti Pietikainen, and David Harwood. Performance evaluation of texture measures with classification based on kullback discrimination of distributions. In *Proceedings of 12th international conference on pattern recognition*, volume 1, pages 582–585. IEEE, 1994.

[90] R. Pappu. *Physical One-Way Functions*. PhD thesis, Massachusetts Institute of Technology, 2001.

[91] Ravikanth Pappu, Ben Recht, Jason Taylor, and Neil Gershenfeld. Physical one-way functions. *Science*, 297(5589):2026–2030, 2002.

[92] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'brien. A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, 5(1):4213, 2014.

[93] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.

[94] Tomaž Prosen, Thomas H Seligman, and Marko Žnidarič. Theory of quantum loschmidt echoes. *Progress of Theoretical Physics Supplement*, 150:200–228, 2003.

[95] Shunfei Qiang, Ke Yuan, Yanyan Cheng, Guoqiang Long, Wenkai Zhang, Xiaofeng Lin, Xiuli Chai, Xiaomin Fang, and Tao Ding. A multicolor carbon dot doped nanofibrous membrane for unclonable anti-counterfeiting and data encryption. *Journal of Materials Chemistry C*, 11(21):7076–7087, 2023.

[96] Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, volume 4, pages 547–562. University of California Press, 1961.

[97] Jonny Roberts, Ibrahim Ethem Bagci, MAM Zawawi, J Sexton, N Hulbert, YJ Noori, MP Young, CS Woodhead, Mohammed Missous, MA Migliorato, et al. Using quantum confinement to uniquely identify devices. *Scientific reports*, 5(1):1–8, 2015.

[98] Ulrich Rührmair, Christian Hilgers, Sebastian Urban, Agnes Weiershäuser, Elias Dinter, Brigitte Forster, and Christian Jirauschek. Optical PUFs reloaded. *Cryptology ePrint Archive*, 2013.

[99] Ulrich Rührmair and Daniel E Holcomb. PUFs at a glance. In *2014 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1–6. IEEE, 2014.

[100] Amit Satpathy, Xudong Jiang, and How-Lung Eng. LBP-based edge-texture features for object recognition. *IEEE Transactions on image Processing*, 23(5):1953–1964, 2014.

[101] Alireza Shamsoshoara, Ashwija Korenda, Fatemeh Afghah, and Sherali Zeadally. A survey on physical unclonable function (PUF)-based security solutions for internet of things. *Computer Networks*, 183:107593, 2020.

[102] Kunal Sharma, Sumeet Khatri, Marco Cerezo, and Patrick J Coles. Noise resilience of variational quantum compiling. *New Journal of Physics*, 22(4):043006, 2020.

[103] Daniel R Simon. On the power of quantum computation. *SIAM journal on computing*, 26(5):1474–1483, 1997.

[104] Boris Škorić. Quantum readout of physical unclonable functions. *International Journal of Quantum Information*, 10(01):1250001, 2012.

[105] W Forrest Stinespring. Positive functions on c*-algebras. *Proceedings of the American Mathematical Society*, 6(2):211–216, 1955.

[106] G Edward Suh and Srinivas Devadas. Physical unclonable functions for device authentication and secret key generation. In *Proceedings of the 44th annual design automation conference*, pages 9–14, 2007.

[107] Giacomo Torlai, Guglielmo Mazzola, Juan Carrasquilla, Matthias Troyer, Roger Melko, and Giuseppe Carleo. Neural-network quantum state tomography. *Nature Physics*, 14(5):447–450, 2018.

[108] Constantino Tsallis. Possible generalization of boltzmann-gibbs statistics. *Journal of statistical physics*, 52:479–487, 1988.

[109] Robert R Tucci. A rudimentary quantum compiler. *arXiv preprint quant-ph/9805015*, 1998.

[110] Armin Uhlmann. The "transition probability" in the state space of a-algebra. *Reports on Mathematical Physics*, 9(2):273–279, 1976.

[111] AV Uvarov and Jacob D Biamonte. On barren plateaus and cost function locality in variational quantum algorithms. *Journal of Physics A: Mathematical and Theoretical*, 54(24):245301, 2021.

[112] Davide Venturelli, Minh Do, Eleanor Rieffel, and Jeremy Frank. Compiling quantum circuits to realistic hardware architectures using temporal planners. *Quantum Science and Technology*, 3(2):025004, 2018.

[113] John VonNeumann. Mathematische grundlagen der quantenmechanik. 1932.

[114] Antian Wang, Weihang Tan, Yuejiang Wen, and Yingjie Lao. NoPUF: A novel PUF design framework toward modeling attack resistant PUFs. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 68(6):2508–2521, 2021.

[115] Chau-Wai Wong and Min Wu. Counterfeit detection using paper PUF and mobile cameras. In *2015 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6. IEEE, 2015.

[116] Bo Yang and Songcan Chen. A comparative study on local binary pattern (LBP) based face recognition: LBP histogram versus LBP image. *Neurocomputing*, 120:365–379, 2013.

[117] Henry Yuen. An improved sample complexity lower bound for (fidelity) quantum state tomography. *Quantum*, 7:890, 2023.

[118] Tongtong Zhang, Lingzhi Wang, Jing Wang, Zhongqiang Wang, Madhav Gupta, Xuyun Guo, Ye Zhu, Yau Chuen Yiu, Tony KC Hui, Yan Zhou, et al. Multimodal dynamic and unclonable anti-counterfeiting using robust diamond microparticles on heterogeneous substrate. *Nature Communications*, 14(1):2507, 2023.

[119] Yuan-Yuan Zhao, Zhibo Hou, Guo-Yong Xiang, Yong-Jian Han, Chuan-Feng Li, and Guang-Can Guo. Experimental demonstration of efficient quantum state tomography of matrix product states. *Optics Express*, 25(8):9010–9018, 2017.